

A CONDITION NUMBER FOR THE TENSOR RANK DECOMPOSITION

NICK VANNIEUWENHOVEN

ABSTRACT. The tensor rank decomposition problem consists of recovering the unique set of parameters representing a robustly identifiable low-rank tensor when the coordinate representation of the tensor is presented as input. A condition number for this problem measuring the sensitivity of the parameters to an infinitesimal change to the tensor is introduced and analyzed. It is demonstrated that the absolute condition number coincides with the inverse of the least singular value of Terracini’s matrix. Several basic properties of this condition number are investigated.

1. INTRODUCTION

A tensor $\mathfrak{A} \in \mathbb{F}^{n_1} \otimes \cdots \otimes \mathbb{F}^{n_d}$, where $\mathbb{F} = \mathbb{C}$ or \mathbb{R} , and $n_k \geq 2$, $k = 1, \dots, d$, lives in the tensor product of vector spaces. It can be represented as a d -array after choosing bases. If the coordinate representation of \mathfrak{A} is given with respect to the standard tensor basis, then we can explicitly write

$$\mathfrak{A} = \sum_{i_1=1}^{n_1} \sum_{i_2=1}^{n_2} \cdots \sum_{i_d=1}^{n_d} a_{i_1, i_2, \dots, i_d} \mathbf{e}_{i_1}^1 \otimes \mathbf{e}_{i_2}^2 \otimes \cdots \otimes \mathbf{e}_{i_d}^d,$$

where \mathbf{e}_i^k is the i th standard basis vector of \mathbb{F}^{n_k} . Tensors of the form $\mathbf{a}^1 \otimes \mathbf{a}^2 \otimes \cdots \otimes \mathbf{a}^d$ are called rank-1 tensors. As in the 2-factor case, i.e., matrices, the above form is, in general, not the minimal way of writing a tensor as a linear combination of rank-1 tensors. This leads to a different type of tensor decomposition:

$$(1) \quad \mathfrak{A} = \sum_{i=1}^r \mathbf{a}_i^1 \otimes \mathbf{a}_i^2 \otimes \cdots \otimes \mathbf{a}_i^d, \quad \text{where } \mathbf{a}_i^k \in \mathbb{F}^{n_k}.$$

If the number of terms r is minimal, then this decomposition called a *tensor rank decomposition* [47, 48]. The number r is then called the *rank* of the tensor. In the literature, decomposition (1) is sometimes also called Candecomp [16], Parafac [43], Canonical Polyadic decomposition, or CP decomposition. Here, we will call (1), with r not necessarily minimal, an r -term decomposition, and if r is minimal, then we call it a rank- r decomposition. One key property of the latter—but not of the former—is that a rank- r tensor with $d \geq 3$ often admits an *essentially unique* decomposition [26, 38, 53]: the rank-1 tensors appearing in decomposition (1) are uniquely determined. For instance, the well-known Kruskal criterion [53] states that (1) is the essentially unique decomposition if

$$r \leq \frac{1}{2}(k_1 + k_2 + k_3 - 2),$$

where k_i is the *Kruskal rank* of $A_i = \{\mathbf{a}_j^i\}_{j=1}^r$, i.e., the integer k_i such that every subset of A_i of cardinality k_i forms a linearly independent set.

1.1. Motivation. The essential uniqueness of a tensor rank decomposition turns out to be very useful in several applications, such as in the parameter identification problem in latent variable models [3, 4]. Models whose parameters may be inferred with this decomposition (and symmetrical variants) were recently surveyed in a unified tensor-based framework by Anandkumar, Ge, Hsu, Kakade, and Telgarsky [4]; they include exchangeable single topic models, hidden Markov models, and Gaussian mixture models. The blind source separation problem in signal processing is another example where the uniqueness of

2000 *Mathematics Subject Classification.* 65F35, 15A69, 15A72, 65H04, 14Q15, 14N05, 58A05, 58A07.

Key words and phrases. tensor rank decomposition, condition number, stability analysis, Terracini’s matrix.

The author’s research was supported by a Postdoctoral Fellowship of the Research Foundation–Flanders (FWO).

the (symmetric) tensor rank decomposition of a cumulant tensor is key in identifying the underlying statistically independent signals [28]. A crucial application of the essential uniqueness of the tensor rank decomposition is found in the chemometrics literature, in particular in fluorescence spectroscopy where its use was introduced by Appelhof and Davidson [5]. Since I will treat this application as a prototypical use of the tensor rank decomposition in parameter identification problems—such as all aforementioned applications—I recall the discussion in [74, Section 2.3]. For a sufficiently diluted solution, some type of molecules called fluorophores will absorb a certain amount of light at the i th wavelength and re-emit a fraction of the absorbed energy as light at the j th wavelength. Mathematically, $x_{i,j} = \chi \lambda_i \mu_j$, where $x_{i,j}$ is the intensity of the light emitted at the j th wavelength when the fluorophore is excited by light at the i th wavelength, λ_i models the fraction of light absorbed at the i th wavelength, μ_j models the fraction of light emitted at the j th wavelength, and χ is some chemical constant proportional to the concentration of the fluorophore in the solution. When r fluorophores occur in a solution and the interactions among them are negligible, then their signal is additive: $x_{i,j} \approx \sum_{\ell=1}^r \chi_\ell \lambda_{i,\ell} \mu_{j,\ell}$. If several of such mixtures with varying concentrations of fluorophores are simultaneously analyzed, the model for the intensity of the emitted light is a tensor rank decomposition: $x_{i,j,k} \approx \sum_{\ell=1}^r \chi_{k,\ell} \lambda_{i,\ell} \mu_{j,\ell}$; the model is only approximately valid, as it rests on some assumptions such as a sufficiently diluted solution and no interactions among the fluorophores. As stated in [74, Section 2.3], “the uniqueness property of [tensor rank decompositions] guarantees that, if [the above model] is adequate, [...] the estimated loading vectors for the ℓ th component are readily interpretable as concentration profiles $[\chi_\ell]$ and emission and excitation spectra ($[\mu_\ell]$ and $[\lambda_\ell]$ respectively) of the ℓ th fluorophore.” This allows for an identification of the fluorophores present in a chemical mixture of unknown composition. While the tensor \mathfrak{X} may be obtained from physical measurements, such a process is always subject to measurement errors because of the finite precision of the measuring equipment. Consequently, only an approximate low-rank decomposition may be anticipated.

The foregoing application highlights that one rarely works with the true low-rank tensor, however one still desires to interpret the uniquely defined rank-1 tensors appearing in the true decomposition. One desires knowing, or at least estimating, how much the individual rank-1 terms in the decomposition computed from the approximate data differ from the true underlying rank-1 tensors. Such an analysis will be the topic of the paper. In order to state the main question more precisely, one needs to specify more carefully which problem we are trying to solve. In the literature, the parameters of an r -term decomposition are often succinctly represented using *factor matrices*. For the present purpose, dealing with a particular vectorization of the factor matrices is more convenient. Define the bijections

$$\begin{aligned} \text{vec} : \mathbb{F}^{n_1} \times \cdots \times \mathbb{F}^{n_d} &\rightarrow \mathbb{F}^{\Sigma+d} & \text{vecr} : (\mathbb{F}^{n_1} \times \cdots \times \mathbb{F}^{n_d})^{\times r} &\rightarrow \mathbb{F}^{r(\Sigma+d)} \\ (\mathbf{a}_i^1, \mathbf{a}_i^2, \dots, \mathbf{a}_i^d) &\mapsto \begin{bmatrix} \mathbf{a}_i^1 \\ \mathbf{a}_i^2 \\ \vdots \\ \mathbf{a}_i^d \end{bmatrix}, \quad \text{and} & (b_1, b_2, \dots, b_r) &\mapsto \begin{bmatrix} \text{vec}(b_1) \\ \text{vec}(b_2) \\ \vdots \\ \text{vec}(b_r) \end{bmatrix}, \end{aligned}$$

where $\Sigma = \sum_{k=1}^d (n_k - 1)$. Let $\Pi = \prod_{k=1}^d n_k$. With these definitions in place, we can introduce the “tensor computation” function

$$\begin{aligned} f : \mathbb{F}^{r(\Sigma+d)} &\rightarrow \sigma_r^0(\mathcal{S}_{\mathbb{F}}) \subset \mathbb{F}^{\Pi} \\ \mathbf{p} = \text{vecr}(b_1, b_2, \dots, b_r) &\mapsto \sum_{i=1}^r \mathbf{a}_i^1 \otimes \mathbf{a}_i^2 \otimes \cdots \otimes \mathbf{a}_i^d, \end{aligned}$$

which takes the *vectorized factor matrices* \mathbf{p} as input and sends them to the tensor they represent. Now we can formally state that the *tensor rank decomposition problem* with as input the rank- r tensor $\mathfrak{A} \in \mathbb{F}^{n_1} \otimes \cdots \otimes \mathbb{F}^{n_d} \cong \mathbb{F}^{\Pi}$ consists of finding a vector $\mathbf{x} \in \mathbb{F}^{r(\Sigma+d)}$ in the preimage of f at \mathfrak{A} , so that $\mathfrak{A} = f(\mathbf{x})$. We adopt the convenient notation

$$f^\dagger(\mathfrak{A}) = \{\mathbf{x} \in \mathbb{F}^{r(\Sigma+d)} \mid f(\mathbf{x}) = \mathfrak{A}\}$$

for denoting the fiber of f at \mathfrak{A} . Since a tensor rank decomposition is not unique in the strict sense, one should be wary that f is *never* an injective function. Hence, f^\dagger is always a *set-valued map* in the sense of Dontchev and Rockafellar [39]. If \mathfrak{A} does not admit an exact rank- r decomposition, then $f^\dagger(\mathfrak{A}) = \emptyset$.

The goal of this paper is understanding the relationship between the essentially unique decompositions $f^\dagger(\mathfrak{A})$ and $f^\dagger(\mathfrak{A}')$ when $\|\mathfrak{A} - \mathfrak{A}'\| \leq \epsilon$ and both \mathfrak{A} and \mathfrak{A}' are r -identifiable tensors through a structured perturbation analysis of the inverse problem $\mathfrak{A} = f(\mathbf{x})$. Specifically, I will derive a *condition number* of the tensor rank decomposition problem that captures the first-order behavior of aforementioned relationship.

A condition number of a function $g : \mathbb{F}^m \rightarrow \mathbb{F}^n$ is an elementary, well-entrenched property in numerical analysis [15, 46, 69, 81] that measures the sensitivity of g to infinitesimal changes in the input. In this paper, I consider the *structured* condition number [15], which differs only from the usual definition in restricting the domain of g to some connected domain $\mathcal{D} \subset \mathbb{F}^m$; for instance, \mathcal{D} could be the set of rank- r tensors. Formally, the structured condition number measures the maximum relative change at the output of the function when infinitesimally perturbing the input $\mathbf{x} \in \mathcal{D}$ to another $\mathbf{x} + \Delta\mathbf{x} \in \mathcal{D}$:

$$(2) \quad \kappa_{\alpha,\beta}(\mathbf{x}) = \lim_{\epsilon \rightarrow 0} \max_{\substack{\|\Delta\mathbf{x}\| \leq \epsilon, \\ \mathbf{x} + \Delta\mathbf{x} \in \mathcal{D}}} \frac{\|g(\mathbf{x} + \Delta\mathbf{x}) - g(\mathbf{x})\|_\alpha / \|g(\mathbf{x})\|_\alpha}{\|\Delta\mathbf{x}\|_\beta / \|\mathbf{x}\|_\beta},$$

where $\|\cdot\|_\alpha$ and $\|\cdot\|_\beta$ are norms on \mathbb{F}^n and \mathbb{F}^m respectively, and $\mathbf{x} \in \mathcal{D} \subset \mathbb{F}^m$. If $\mathcal{D} = \mathbb{F}^m$ then the usual condition number is obtained.

The main application of condition numbers consists of obtaining an asymptotically sharp *a posteriori* bound on the relative forward error, by grace of the observation

$$(3) \quad \frac{\|g(\mathbf{x} + \Delta\mathbf{x}) - g(\mathbf{x})\|_\alpha}{\|g(\mathbf{x})\|_\alpha} \lesssim \kappa_{\alpha,\beta}(\mathbf{x}) \cdot \frac{\|\Delta\mathbf{x}\|_\beta}{\|\mathbf{x}\|_\beta},$$

where $\mathbf{x}, \mathbf{x} + \Delta\mathbf{x} \in \mathcal{D}$ and $\|\Delta\mathbf{x}\|_\beta$ should be sufficiently close to zero. That is, *the relative forward error is asymptotically bounded by the product of the relative backward error multiplied with the condition number*, which is a well-known rule of thumb in numerical analysis. This observation is particularly important in applications, because the forward error usually cannot be computed for a practical problem, as you typically would have to know the true solution $g(\mathbf{x} + \Delta\mathbf{x})$; here \mathbf{x} would represent the observed input to problem g which is typically corrupted in varying degrees by measurement, modeling, computation and representation errors, whereas $\mathbf{x} + \Delta\mathbf{x}$ would represent the true input. The key advantage of a computable condition number of a problem g is that the forward error can be estimated from *a posteriori* information, namely the backward error and the condition number. For example, in the aforementioned chemometrics application, the underlying physical model ensures that the theoretical tensor \mathfrak{A} admits a unique rank decomposition of low rank as in (1), say, with vectorized factor matrices $\mathbf{x} + \Delta\mathbf{x}$: $\mathfrak{A} = f(\mathbf{x} + \Delta\mathbf{x})$. Each of the rank-1 terms uniquely corresponds with an individual chemical compound in the fluorescent mixture. In practice, the tensor $\tilde{\mathfrak{A}}$ that one obtained by measuring the reflected light in an experiment contains measurement errors. This may be regarded as a perturbation of \mathfrak{A} of small Frobenius norm, say, $\|\mathfrak{A} - \tilde{\mathfrak{A}}\|_F \leq \mu$. A low-rank approximate decomposition of $\tilde{\mathfrak{A}}$ is then constructed using a numerical method for computing an r -term decomposition. This results in an approximation $\mathfrak{A} \approx \tilde{\mathfrak{A}} \approx \hat{\mathfrak{A}} = f(\mathbf{x})$ with factor matrices \mathbf{x} . If $\|\hat{\mathfrak{A}} - \tilde{\mathfrak{A}}\|_F = \nu$, then we know that the distance between the approximate decomposition $\hat{\mathfrak{A}} = f(\mathbf{x})$ computed from the sampled data $\tilde{\mathfrak{A}}$ and the true tensor corresponding to the physical phenomenon \mathfrak{A} is at most $\mu + \nu$. If we compute the condition number κ of the tensor rank decomposition $\hat{\mathfrak{A}} = f(\mathbf{x})$, then it reveals an upper bound on the relative difference between the computed approximate factor matrices \mathbf{x} and the true factor matrices $\mathbf{x} + \Delta\mathbf{x}$. If this relative forward error is small, then one can safely match the approximate rank-1 terms to one or more chemical compounds, as the true excitation spectra will have been well approximated. In several data-analytic applications where the factor matrices themselves are analyzed and linked to certain phenomena or interpretations, all too often the analysis is stopped when a small backward error $\mu + \nu$ is obtained without considering the well-known fact in numerical analysis that a small backward error must not entail a small forward error. A small backward error $\|\mathfrak{A} - \hat{\mathfrak{A}}\|_F \leq \mu + \nu$ entails a small forward error $\|\Delta\mathbf{x}\|$ only if the problem is *well-conditioned*, i.e., if the condition number is small.

1.2. Contributions. The main contribution of this paper is a definition of a condition number for the rank decomposition problem f^\dagger . This condition number is well-defined for every rank decomposition that satisfies a technical condition called robust r -identifiability. This condition is generally satisfied, such as for a generic¹ rank- r decomposition in every tensor space that is generically r -identifiable—see Lemma 3; most spaces are conjectured to be of this kind, see, e.g., [12, 26]. The prime result, Theorem 7, may be stated informally as follows.

Theorem 1. *Let $N = r(\Sigma + 1) < \Pi$, $N' = \min\{r(\Sigma + d), \Pi\}$, and let $\mathfrak{A} = f(\mathbf{x})$ be a generic rank- r tensor in $\mathbb{F}^{n_1 \times n_2 \times \dots \times n_d}$. If $J = [\partial f_i / \partial x_j]_{i,j}$ denotes the Jacobian matrix of f evaluated at \mathbf{x} and $\varsigma_1 \geq \varsigma_2 \geq \dots \geq \varsigma_N \geq \varsigma_{N+1} = \dots = \varsigma_{N'} = 0$ are its singular values, then $\kappa_A = \varsigma_N^{-1}$ is an absolute condition number of the tensor rank decomposition problem $\mathfrak{A} = f(\mathbf{x})$.*

The proof of this theorem is rather technical, and it should become clear to the reader that several of the technical issues could have been overcome if a slightly different definition of a condition number was investigated—in particular, the reader may notice that Part I of the proof of Theorem 7 contains in it the nuggets of a definition of a (set of) alternative condition numbers. The technical issues are mathematically easy to overcome by reducing the number of parameters in the vectorized factor matrices to the (expected) dimension (as algebraic variety) of the set of rank- r tensors. This amounts to removing the so-called *scaling indeterminacies* that are inherent in the definition of f , resulting in a Jacobian matrix with linearly independent columns. The scaling indeterminacies can be removed in several ways, perhaps most straightforwardly by simply removing some of the variables (as in Part I of the proof of Theorem 7). However, I stress that this approach leads to a condition number of a *different* problem \hat{f}^\dagger , simply because the function \hat{f} mapping the (minimal) set of parameters to a tensor is different from f . This approach introduces some challenges of its own, mainly related to the fact that each choice of elimination of variables leads to a different tensor rank decomposition problem, and hence to a different condition number. While it leads to very interesting questions such as determining which elimination of variables one should use to minimize the condition number in a given problem instance, this paper provides an answer to Nicaise’s question:²

Does f^\dagger admit an intrinsic condition number, i.e., one that does not depend (even implicitly) on the choice of affine chart (or, equivalently, on the choice of elimination of variables)?

This is the question that I believe Theorem 7 answers. I will not investigate the alternative set-of-condition-numbers approach in this paper.

Another way of looking at Theorem 1 is that it provides an answer to the following research question:

Does the singular value ς_N of the Jacobian of f admit any kind of geometrical interpretation?

This Jacobian matrix lies at the heart of gradient-based optimization algorithms for the computation of (approximate) tensor rank decompositions, such as the algorithms in [2, 45, 58, 59, 62, 63, 67, 75, 77]. That the inverse of this singular value admits an interpretation as a condition number is somewhat of a free lunch: one needs to construct the Jacobian matrix (or its Gram matrix) in every step of the optimization algorithm anyway, and at the final iteration one obtains an estimate of the conditioning of the computed decomposition with little extra effort. In addition, it was already observed in the chemometrics literature that this singular value seems to be related to the stability of the tensor rank decomposition. According to Tomasi [74], “several diagnostic parameters have been used in the chemometrics literature to measure the difficulty of [tensor rank decomposition] problems, but many seem to fall short of giving consistent and concise information.” Tomasi proceeds by listing the various criteria that have been proposed, among them “the Jacobian’s condition number and singular values” that were studied in [60, 61, 75, 76]. This paper proves that the inverse of the least singular value of the Jacobian matrix has a clear, consistent and concise interpretation: *it is a condition number* of the tensor rank decomposition problem.

¹A property that may be admitted by the elements of a set S is called generic if all elements of $S \setminus N$ admit the property, where $N \subset S$ has measure zero in S in some topology. Note that the notion of generality depends on the underlying topology. A property may, hence, be generic with respect to one topology, while it is not generic in another.

²Remark that Nicaise’s question explicitly excludes taking the minimum or maximum over all possible modified tensor rank decomposition problems \hat{f}^\dagger .

1.3. Related work. Despite the importance of a perturbation theory for the tensor rank decomposition problem, heretofore few attempts have been undertaken to develop such a theory.

As mentioned before, [74, p. 31] contains an overview of several criteria that have been proposed in the chemometrics literature for assessing the difficulty of the tensor rank decomposition problem. However, none of the cited references derives a condition number as such. It will be shown in Theorem 7 that the least singular value of the Jacobian matrix that was already proposed as a heuristic for checking stability is, in fact, a condition number, and, as such, admits the useful interpretation in (3).

Robust identifiability of the parameters of a tensor rank decomposition was investigated by Bhaskara, Charikar, and Vijayaraghavan [7]. It is a necessary technical requirement for the definition of the condition number. In addition, it is shown in [7] how these robustness results relate to the accuracy with which the parameters of a tensor rank decompositions, i.e., the factor matrices, can be identified.

Perhaps the result closest in spirit to this work are the Cramér–Rao bounds (CRB) for the tensor rank decomposition problem that have been investigated in [52, 57, 72, 73]. This CRB measures the stability of a tensor rank decomposition in a statistical framework, wherein additive Gaussian noise is assumed to corrupt the factor matrices. The quantity of interest is the squared angular error between the true parameters and the estimated parameters [73]. The Cramér–Rao lower bound (CRLB) for estimating the parameters \mathbf{p} of the tensor rank decomposition of $\mathfrak{A} = f(\mathbf{p})$ is then defined [73] as the inverse of the Fisher information matrix $\frac{1}{\sigma^2} J^T J$, where J is as in Theorem 1 and σ^2 is the variance of the Gaussian noise. However, this matrix is not invertible, so in [73] it is suggested to alter the tensor rank decomposition problem by removing some of the parameters. This implies that the CRLB depends on the particular elimination of the variables that is chosen; hence, it is not an *intrinsic* measure of stability of the tensor rank decomposition problem as defined in this paper.

As an intermezzo, [73] states as an open problem whether “stability of the [tensor rank decomposition] problem implies its essential uniqueness.” The answer is an unqualified no. There are explicit counterexamples where the generic tensor is not r -identifiable, yet whereby the N th singular value of the Jacobian J is nonzero, so that the tensor still has finitely many decompositions—for a conjecturally complete list of such tensors, see [26, Theorem 1.1]. In these cases, the CRLB will be finite, but the tensor is not identifiable. Counterexamples are not limited to those appearing in [26, Theorem 1.1], however they expected to be intrinsically tied up with singular points of generically r -identifiable Segre varieties, as Example 4.2 in [26] shows.

1.4. Outline. The rest of this paper is structured as follows. In the next section, some necessary preliminaries from (semi-)algebraic geometry are recalled. The main object of study, Terracini’s matrix, is introduced. Section 3 analyzes the fiber f^\dagger . It is shown that the fiber is usually unique up to the action of some multiplicative group \mathcal{T} . The concept of robust r -identifiability is introduced as a weak technical assumption. Essentially it allows us to assume that a neighborhood of a point on the r -secant variety of the Segre variety exists wherein all points are smooth, of maximal dimension, and r -identifiable. Section 4 introduces a distance measure that ignores the action of the multiplicative group \mathcal{T} acting on f^\dagger . Terracini’s matrix has a nontrivial kernel originating from \mathcal{T} , however Section 5 proves that it may be ignored completely in an analysis of the condition number with respect to the proposed distance measure. This lemma is proved by constructing a kind of approximate Newton-type algorithm that replaces a vector in the kernel of Terracini’s matrix with a vector in the orthogonal complement of the kernel of much smaller norm. Thereafter, the fact that the inverse of the least nonzero singular value of Terracini’s matrix corresponds with an absolute condition number with respect to the distance measure of Section 4 is proved in Section 6. Section 7 continues by defining a particular condition number and presents an algorithm for computing it; the basic properties of this condition number are further analyzed in Section 8. Numerical experiments illustrating the behavior of the proposed condition number are presented in Section 9. A discussion of the conclusions and future research directions concludes the paper in Section 10.

Acknowledgements. This paper benefited immeasurably from the fruitful discussions I had with B. Jeuris, K. Meerbergen, and G. Ottaviani. I am particularly indebted to J. Nicaise for inquiring about an intrinsic definition of the condition number. I thank P. Breiding for detailed comments on an earlier version of this manuscript.

2. PRELIMINARIES

Several requisite results concerning the tensor rank decomposition are recalled with a particular focus on some basic concepts from algebraic geometry. While the varieties that we study admit the structure of projective varieties, we will present the results in affine space. The reason is that in the applications where we envisage that a condition number would be of interest, one commonly works in affine space. In addition, the definition of a condition number in this paper is inherently local, so that we may always restrict ourselves to affine varieties. Before proceeding, some notation is fixed.

Notation and conventions. Varieties are typeset in a calligraphic uppercase font (\mathcal{S}, \mathcal{V}), tensors in an uppercase fraktur font (\mathfrak{A}), matrices in uppercase letters (A, U, V), vectors in boldface lowercase letters (\mathbf{a}, \mathbf{x}), and scalars, as well as points on varieties, in lowercase letters (a, b, λ, p, q). The symbol \mathbb{F} denotes either the real field \mathbb{R} or the complex field \mathbb{C} . For a matrix $A \in \mathbb{F}^{m \times n}$, A^T is its transpose and A^H is its conjugate transpose. The Euclidean norm of $\mathbf{v} \in \mathbb{F}^m$ is $\|\mathbf{v}\| = \sqrt{\mathbf{v}^H \mathbf{v}}$. A block diagonal matrix with diagonal blocks A_1, \dots, A_d is denoted by $\text{diag}(A_1, \dots, A_d)$. By $\text{span}(A)$ the column span of the matrix A is meant.

The prototypical tensor in this paper is the rank- r tensor $\mathfrak{A} \in \mathbb{F}^{n_1 \times n_2 \times \dots \times n_d}$. Hence, the scalar d is used exclusively for the order of \mathfrak{A} , r denotes its rank, the scalars n_1, n_2, \dots, n_d are reserved for \mathfrak{A} 's dimensions, and we also define

$$\Sigma = \sum_{k=1}^d (n_k - 1), \quad N = r(\Sigma + 1), \quad \text{and} \quad \Pi = \prod_{k=1}^d n_k,$$

which are respectively the dimension of the variety of rank-1 tensors embedded in $F = \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \dots \otimes \mathbb{F}^{n_d}$, the expected dimension of the smallest variety containing the tensors of rank r in F , and the dimension of F .

2.1. Segre varieties and secants. The Segre variety over $\mathbb{F} = \mathbb{C}$ or \mathbb{R} , denoted by

$$\mathcal{S}_{\mathbb{F}} = \text{Seg}(\mathbb{F}^{n_1} \times \mathbb{F}^{n_2} \times \dots \times \mathbb{F}^{n_d}) \subset \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \dots \otimes \mathbb{F}^{n_d} \cong \mathbb{F}^{\Pi},$$

is the smooth, determinantal variety of rank-1 tensors. Its elements can be parameterized explicitly via the Segre map:

$$\begin{aligned} \text{Seg} : \mathbb{F}^{n_1} \times \mathbb{F}^{n_2} \times \dots \times \mathbb{F}^{n_d} &\rightarrow \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \dots \otimes \mathbb{F}^{n_d} \cong \mathbb{F}^{\Pi} \\ (\mathbf{a}^1, \mathbf{a}^2, \dots, \mathbf{a}^d) &\mapsto \mathbf{a}^1 \otimes \mathbf{a}^2 \otimes \dots \otimes \mathbf{a}^d \end{aligned}$$

where \otimes is the tensor product. For notational brevity, I will consider the Segre map that embeds into \mathbb{F}^{Π} ; hence, in coordinates the tensor product may be realized as the Kronecker product. Note that the 2-factor Segre variety is the set of rank-1 matrices. In this paper, it is often required to refer to the specific vectors $b_i = (\mathbf{a}_i^1, \mathbf{a}_i^2, \dots, \mathbf{a}_i^d)$ in the representation of a point $p_i = \text{Seg}(b_i) \in \mathcal{S}_{\mathbb{F}}$ on a Segre variety $\mathcal{S}_{\mathbb{F}}$. We will call b_i the *representative* of p_i .

Let $r < \frac{\Pi}{\Sigma+1}$ be strictly *subgeneric*, and define the set of tensors of \mathbb{F} -rank at most r as

$$\sigma_r^0(\mathcal{S}_{\mathbb{F}}) = \{p_1 + p_2 + \dots + p_r \mid p_i \in \mathcal{S}_{\mathbb{F}}\} = \{\mathfrak{A} \in \mathbb{F}^{n_1} \otimes \dots \otimes \mathbb{F}^{n_d} \mid \text{rank}_{\mathbb{F}}(\mathfrak{A}) \leq r\}.$$

Over the complex numbers, $\mathbb{F} = \mathbb{C}$, the closures of $\sigma_r^0(\mathcal{S}_{\mathbb{C}})$ in both the Euclidean and Zariski topology agree [54, Corollary 5.1.1.5], so that rank- r tensors form a Euclidean-dense subset of the algebraic variety $\sigma_r(\mathcal{S}_{\mathbb{C}}) = \overline{\sigma_r^0(\mathcal{S}_{\mathbb{C}})}$, where the overline indicates the Zariski-closure. This projective variety is called the *r-secant variety* of the Segre variety $\mathcal{S}_{\mathbb{C}}$. Over the real numbers, $\mathbb{F} = \mathbb{R}$, we could define the *r-secant variety* analogously by taking the closure in the Zariski topology, resulting in an algebraic variety over \mathbb{R} . Unfortunately, $\sigma_r^0(\mathcal{S}_{\mathbb{R}})$ is not Euclidean-dense in $\sigma_r^0(\mathcal{S}_{\mathbb{R}})$ in general. This can be seen, for example, in the occurrence of multiple *typical real ranks* [6, 10, 30, 54], i.e., the fact that there may be several open sets in $\mathbb{R}^{n_1 \times \dots \times n_d}$ each with constant and different \mathbb{R} -rank. For this reason, it is more sensible to consider the closure of $\sigma_r^0(\mathcal{S}_{\mathbb{R}})$ in the Euclidean topology, which results in a semi-algebraic set that we denote by $\sigma_r(\mathcal{S}_{\mathbb{R}})$, see, e.g., [33, Theorem 6.2]. Recall that a semi-algebraic set is the solution set of a system of polynomial equations and inequalities over \mathbb{R} —see [13].

The foregoing showed that for both $\mathbb{F} = \mathbb{C}$ and \mathbb{R} the tensors of \mathbb{F} -rank equal to r form a Euclidean-dense subset of $\sigma_r(\mathcal{S}_{\mathbb{F}})$. Taking the closure of $\sigma_r^0(\mathcal{S}_{\mathbb{F}})$ in the Euclidean topology generally introduces new points, i.e., $\mathcal{R} = \sigma_r(\mathcal{S}_{\mathbb{F}}) \setminus \sigma_r^0(\mathcal{S}_{\mathbb{F}}) \neq \emptyset$. This residual set \mathcal{R} contains tensors of rank strictly larger than r , which can be approximated arbitrarily well by tensors of rank r . Such tensors are called border tensors [8, 9]. They cause problems when trying to approximate them by a rank- r tensor, which will be illustrated in Section 9.5.

2.2. Dimension. A fundamental property of manifolds, varieties and semi-algebraic sets is that at most points they locally resemble a linear space. One way to describe this local linear approximation to a manifold, variety or semi-algebraic set $\mathcal{V} \subset \mathbb{F}^{\Pi}$ at a point $p \in \mathcal{V}$ is the *Zariski tangent space*; see respectively [55, p. 54], [41, pp. 175–176] and [13, Section 3.3]. This vector space is defined as the span of all vectors in \mathbb{F}^{Π} that are tangent to \mathcal{V} :

$$T_p \mathcal{V} = \text{span}\left(\left\{\frac{d}{dt}p(t) \mid p(t) \in \mathcal{V}, t \in (0, 1], p(0) = p\right\}\right).$$

The dimension of a manifold or variety $\mathcal{V} \subset \mathbb{F}^{\Pi}$ may be defined geometrically as the number n such that $\dim T_p \mathcal{V} = n$ for all p in an open neighborhood \mathcal{N} of $p_0 \in \mathcal{V}$. For a semi-algebraic set \mathcal{V} the dimension is defined as the dimension of the Zariski-closure $\overline{\mathcal{V}}$ [13, Proposition 2.8.2]. The dimension of the Segre variety $\mathcal{S}_{\mathbb{F}} = \text{Seg}(\mathbb{F}^{n_1} \times \cdots \times \mathbb{F}^{n_d}) \subset \mathbb{F}^{\Pi}$ is well-known [42, 54]:

$$\dim \mathcal{S}_{\mathbb{F}} = \Sigma + 1 = 1 + \sum_{k=1}^d (n_k - 1).$$

The dimension of $\sigma_r(\mathcal{S}_{\mathbb{F}})$ is, in general, not known. Since $\sigma_r^0(\mathcal{S}_{\mathbb{F}})$ is Euclidean-dense in $\sigma_r(\mathcal{S}_{\mathbb{F}})$, an upper bound is readily established by noting that $\sigma_r(\mathcal{S}_{\mathbb{F}})$ is the projection of

$$\overline{\{(p_1, p_2, \dots, p_r), p \mid p = p_1 + p_2 + \dots + p_r, p_1, \dots, p_r \in \mathcal{S}_{\mathbb{F}}\}} \subset (\mathcal{S}_{\mathbb{F}} \times \cdots \times \mathcal{S}_{\mathbb{F}}) \times \mathbb{F}^{\Pi}$$

onto the second factor, where the overline indicates the closure in the Euclidean topology. Hence, $\dim \sigma_r(\mathcal{S}_{\mathbb{F}}) \leq \min\{\Pi, r(\Sigma + 1)\}$, by [13, Proposition 2.8.6] for $\mathbb{F} = \mathbb{R}$ and [42, Theorem 11.12] for $\mathbb{F} = \mathbb{C}$. One expects that equality holds; if this is the case, then we say that the r -secant is *nondefective*. If all secant varieties of a Segre variety $\mathcal{S}_{\mathbb{F}}$ are nondefective, then we say that this Segre variety is nondefective. The problem of determining the dimension of $\sigma_r(\mathcal{S}_{\mathbb{F}})$ has seen much progress during the last two decades, both theoretically [1, 11, 12, 17–21, 23–25, 40] and numerically [26, 31, 80], yet the following is still an open problem:

Conjecture 1 (Abo, Ottaviani, and Peterson [1]). *Let $\mathcal{S}_{\mathbb{F}} = \text{Seg}(\mathbb{F}^{n_1} \times \cdots \times \mathbb{F}^{n_d}) \subset \mathbb{F}^{\Pi}$, be a Segre variety with $d \geq 3$ and $n_1 \geq n_2 \geq \cdots \geq n_d \geq 2$. Then, $\mathcal{S}_{\mathbb{F}}$ is nondefective, unless*

- (1) $(n_1, n_2, n_3) = (4, 4, 3)$,
- (2) $(n_1, n_2, n_3) = (2k + 1, 2k + 1, 3)$ with $k \in \mathbb{N}$,
- (3) $(n_1, n_2, n_3, n_4) = (k + 1, k + 1, 2, 2)$ with $k \in \mathbb{N}$, or
- (4) $n_1 > 1 + \prod_{k=2}^d n_k - \sum_{k=2}^d (n_k - 1)$, i.e., $\mathcal{S}_{\mathbb{F}}$ is unbalanced.

Remark 2. *It is noted that [1] only states the conjecture in the case of $\mathbb{F} = \mathbb{C}$ but the geometrical approach in [26, 80] applies to \mathbb{R} -varieties as well, as it consisted of constructing a subset of the Zariski tangent space at a randomly chosen point of $\sigma_r(\overline{\mathcal{S}_{\mathbb{Z}}})$ and $\sigma_r(\mathcal{S}_{\mathbb{R}})$ in [26] and [80] respectively. By [13, Proposition 2.8.2] the dimension of the \mathbb{R} -variety $\sigma_r(\overline{\mathcal{S}_{\mathbb{R}}})$, where the overline denotes the closure in the Zariski topology, coincides with $\dim \sigma_r(\mathcal{S}_{\mathbb{R}})$.*

2.3. Smooth and regular points. A difficulty in working with varieties—as contrasted to smooth manifolds—is that the local dimension of a complex variety \mathcal{V} is not upper semicontinuous: there may exist points $p \in \mathcal{V}$ where $\dim T_p \mathcal{V} > \dim \mathcal{V}$. A point $p \in \mathcal{V}$ where $\dim T_p \mathcal{V} = \dim \mathcal{V}$ is called a *smooth point* of \mathcal{V} ; otherwise, it is called a *singular point*. The set of singular points is called its *singular locus*, which is a subvariety of dimension strictly less than n . Similarly, the set of smooth points of \mathcal{V} forms an n -dimensional submanifold of \mathbb{C}^{Π} ; it is Zariski-dense in \mathcal{V} . For consistency with the real case, which is treated next, I will call a smooth point of a \mathbb{C} -variety a *regular point*.

For semi-algebraic sets $\mathcal{V} \subset \mathbb{R}^\Pi$, defining the smooth locus is more subtle, because \mathcal{V} may consist of several semi-algebraically connected components of different dimensions. Attempting to define the set of smooth points geometrically as the set of points where $\dim T_p \mathcal{V} = \dim \mathcal{V}$ may in general include points that are singular points in the proper, algebraic definition [13, Definition 3.3.9]. I prefer an alternative geometrical definition that suffices for the purpose of this paper. Every semi-algebraic set \mathcal{V} admits a *Nash stratification*, meaning that it can be decomposed as a finite union of disjoint semi-algebraically connected Nash manifolds [13, Section 9.1]: $\mathcal{V} = \bigcup_i \mathcal{M}_i$, where \mathcal{M}_i is a smooth \mathbb{R} -manifold of dimension $d_i = \dim \mathcal{M}_i$ and $\mathcal{M}_i \cap \mathcal{M}_j = \emptyset$ if $i \neq j$. Let $D = \{i \mid \dim \mathcal{M}_i = \dim \mathcal{V}\}$ be the set of all \mathcal{M}_i of maximal dimension; it is nonempty. I will call $\bigcup_{i \in D} \mathcal{M}_i$ the set of *regular* points of the semi-algebraic set \mathcal{V} . All regular points are smooth but in general the converse is false. The regular locus is neither Zariski-dense nor Euclidean-dense, in general.³

2.4. Terracini's Jacobian. The main object of study in Theorem 1 is the set of singular values of the Jacobian matrix of partial derivatives of f evaluated at \mathbf{p} . This Jacobian matrix is well-known [1, 80] and its structure has been analyzed in the context of optimization algorithms for computing tensor rank decompositions [2, 63, 67, 77]. Assume that we are given a set of representatives $b_i = (\mathbf{a}_i^1, \dots, \mathbf{a}_i^d)$ of the points $p_i = \text{Seg}(b_i) \in \mathcal{S}_{\mathbb{F}}$, and consider the tensor rank decomposition

$$\mathfrak{A} = f(\mathbf{p}) = \sum_{i=1}^r p_i = \sum_{i=1}^r \mathbf{a}_i^1 \otimes \mathbf{a}_i^2 \otimes \dots \otimes \mathbf{a}_i^d,$$

where $\mathbf{p} = \text{vecr}(b_1, b_2, \dots, b_r)$. Then, the Jacobian matrix is given explicitly by

$$(4) \quad T_{\mathbf{p}} = [T_1 \quad T_2 \quad \dots \quad T_r], \text{ where } T_i = [I_{n_1} \otimes \mathbf{a}_i^2 \otimes \dots \otimes \mathbf{a}_i^d \quad \dots \quad \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^{d-1} \otimes I_{n_d}].$$

I will refer to $T_{\mathbf{p}}$ as *Terracini's matrix* associated with the vectorized factor matrices \mathbf{p} because it describes that part of the tangent space to $\sigma_r(\mathcal{S}_{\mathbb{F}})$ that is obtained from an application of Terracini's Lemma [54, 71]. At regular points of $\sigma_r(\mathcal{S}_{\mathbb{F}})$, the column span of Terracini's matrix coincides with the Zariski tangent space. It then follows from Conjecture 1 that $T_{\mathbf{p}}$ is expected to be of maximal rank $N = r(\Sigma + 1)$ for subgeneric $r \leq \frac{\Pi}{\Sigma+1}$, modulo the known exceptions presented in Conjecture 1. For $\mathbb{F} = \mathbb{C}$ this implies that Terracini's matrix is of rank N on a dense subset of $\mathbb{F}^{r(\Sigma+d)}$, while for $\mathbb{F} = \mathbb{R}$ we can only conclude that its rank is N on an open subset of $\mathbb{F}^{r(\Sigma+d)}$.

As the number of columns of $T_{\mathbf{p}}$ is $r(\Sigma + d) > N$, there is a kernel of dimension at least $r(d - 1)$. If $\sigma_r(\mathcal{S}_{\mathbb{F}})$ has the expected dimension, then the kernel can be described explicitly as follows. A basis of the kernel of T_{p_i} is

$$(5a) \quad K_i = \left\{ \begin{bmatrix} \mathbf{a}_i^1 \\ -\mathbf{a}_i^2 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} \mathbf{a}_i^1 \\ 0 \\ -\mathbf{a}_i^3 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \dots, \begin{bmatrix} \mathbf{a}_i^1 \\ 0 \\ 0 \\ \vdots \\ 0 \\ -\mathbf{a}_i^d \end{bmatrix} \right\} = \{\mathbf{k}_i^2, \mathbf{k}_i^3, \dots, \mathbf{k}_i^d\},$$

which contains exactly $d - 1$ linearly independent basis vectors. It is straightforward to check that $T_i K_i = 0$. It is an exercise to verify that

$$(5b) \quad K = \{\mathbf{k}_i^2 \otimes \mathbf{e}_i, \mathbf{k}_i^3 \otimes \mathbf{e}_i, \dots, \mathbf{k}_i^d \otimes \mathbf{e}_i\}_{i=1}^r,$$

where \mathbf{e}_i is the i th standard basis vector of \mathbb{F}^r , forms a basis of the kernel of $T_{\mathbf{p}}$.

³With some additional effort the smooth locus could be defined accurately so that it would be a Euclidean-dense subset of \mathcal{V} but for the purpose of this paper the regular locus suffices since the condition number in Theorems 1 and 7 will be ∞ for non-regular points.

3. THE SOLUTION FIBER AND IDENTIFIABILITY

Since f is never an injective function, the fiber $f^\dagger(\mathfrak{A})$ of f at a rank- r tensor \mathfrak{A} contains several elements. Two trivial sources of this multivaluedness are well-known [53]. The *permutation indeterminacy* arises because the order in which the rank-1 terms appear in (1) is not uniquely determined by \mathfrak{A} ; if $\mathbf{p} = \text{vecr}(b_1, b_2, \dots, b_r) \in f^\dagger(\mathfrak{A})$ is a decomposition of \mathfrak{A} , then so is $\mathbf{p}_\pi = \text{vecr}(b_{\pi_1}, b_{\pi_2}, \dots, b_{\pi_r}) \in f^\dagger(\mathfrak{A})$ for every permutation π of $\{1, 2, \dots, r\}$. If a rank-1 tensor is represented in the standard way by a d -tuple of vectors $(\mathbf{a}^1, \dots, \mathbf{a}^d) \in \mathbb{F}^{n_1} \times \dots \times \mathbb{F}^{n_d}$, then this representation suffers from a *scaling indeterminacy* in the sense that $(\alpha_1 \mathbf{a}^1, \dots, \alpha_d \mathbf{a}^d)$ with $\alpha_1 \cdots \alpha_d = 1$ represents the same rank-1 tensor. Two rank decompositions are considered to be *essentially equal* if the rank-1 tensors appearing in their respective decompositions are equal.⁴

Let us be more concrete about the aforementioned relation. Define⁵

$$\begin{aligned} \mathcal{D} &= \{ \text{diag}((\alpha_2 \cdots \alpha_d)^{-1} I_{n_1}, \alpha_2 I_{n_2}, \dots, \alpha_d I_{n_d}) \mid \alpha_2, \dots, \alpha_d \in \mathbb{F} \setminus \{0\} \}, \\ \mathcal{B} &= \{ \text{diag}(D_1, D_2, \dots, D_r) \mid D_1, D_2, \dots, D_r \in \mathcal{D} \}, \\ \mathcal{P} &= \{ P \otimes I_{\Sigma+d} \mid P \text{ is an } r \times r \text{ permutation matrix} \}, \text{ and} \\ \mathcal{T} &= \mathcal{P}\mathcal{B} = \{ PB \mid P \in \mathcal{P} \text{ and } B \in \mathcal{B} \}. \end{aligned}$$

It is not difficult to verify that for every $P \otimes I_{\Sigma+d} \in \mathcal{P}$ and diagonal matrices $D_i \in \mathbb{F}^{\Sigma+d \times \Sigma+d}$ one has

$$(6) \quad (P \otimes I_{\Sigma+d}) \cdot \text{diag}(D_1, D_2, \dots, D_r) \cdot (P \otimes I_{\Sigma+d})^T = \text{diag}(D_{\pi_1}, D_{\pi_2}, \dots, D_{\pi_r}),$$

where π is the permutation represented by P . It is straightforward to check that both \mathcal{B} and \mathcal{P} are multiplicative groups. From thence one can verify that \mathcal{T} satisfies the properties of a group, however I present here a more elegant alternative that was suggested to me by P. Breiding: (6) defines a map $\psi : \mathcal{P} \rightarrow \text{Aut}(\mathcal{B})$ that takes $x \in \mathcal{P}$ to $\psi_x \in \text{Aut}(\mathcal{B})$ defined by $\psi_x : y \mapsto xyx^{-1}$. Hence, it follows immediately (see, e.g., [65, Chapter 7]) that \mathcal{T} is the semidirect product $\mathcal{T} = \mathcal{P} \rtimes \mathcal{B}$, which is a multiplicative group. Writing $\mathbf{p} \sim \mathbf{q}$ if and only if $\mathbf{p} = T\mathbf{q}$ for some $T \in \mathcal{T}$,⁶ it is now easy to verify that \sim is an equivalence relationship and that f is a morphism for \sim . If $\mathbf{p} \sim \mathbf{q}$ then \mathbf{p} and \mathbf{q} represent tensor rank decompositions that are essentially equal, as we defined above. Naturally, if $\mathbf{p} \sim \mathbf{q}$ then $f(\mathbf{p}) = f(\mathbf{q})$, however *the converse is not a priori true*. If the converse is also true, i.e., $\mathfrak{A} = f(\mathbf{p}) = f(\mathbf{q})$ implies that $\mathbf{p} \sim \mathbf{q}$, then we say that \mathfrak{A} is *r -identifiable*. In other words, if

$$|f^\dagger(\mathfrak{A}) / \sim| = 1,$$

then \mathfrak{A} is of rank r and has one essentially unique tensor rank decomposition.

In the case of the tensor rank decomposition, there is very extensive evidence that most tensors of small rank are identifiable, substantially simplifying the interpretation of the condition number that will be explored in this paper. As is well-known, Kruskal [53] proposed an explicit criterion for testing whether $\mathfrak{A} = f(\mathbf{x}) \in \sigma_r^0(\mathcal{S}_{\mathbb{F}})$ is r -identifiable, assuming that \mathbf{x} is provided as input to the test. Several other such criteria have been proposed since—e.g., [26, 35, 50, 66]. More powerful theorems and conjectures are known in the setting of *generic r -identifiability*. A Segre variety $\mathcal{S}_{\mathbb{F}}$ is called generically r -identifiable if there exists a Zariski-closed set $Z \subset \sigma_r(\mathcal{S}_{\mathbb{F}}) \setminus \sigma_r^0(\mathcal{S}_{\mathbb{F}})$ such that every $\mathfrak{A} \in \sigma_r(\mathcal{S}_{\mathbb{F}}) \setminus Z$ is r -identifiable. Several papers have proved generic r -identifiability of the secant varieties $\sigma_r(\mathcal{S}_{\mathbb{F}})$, such as [12, 25, 38, 44, 70]. In the case of the algebraically closed field $\mathbb{F} = \mathbb{C}$, the results are much more developed than in the real case; the following conjecture paints a nearly complete picture:

Conjecture 2 (Chiantini, Ottaviani, and Vannieuwenhoven [26]). *Let $\mathcal{S}_{\mathbb{C}} = \text{Seg}(\mathbb{C}^{n_1} \times \mathbb{C}^{n_2} \times \dots \times \mathbb{C}^{n_d}) \subset \mathbb{C}^{\Pi}$ be a Segre variety with $d \geq 3$ and $n_1 \geq n_2 \geq \dots \geq n_d \geq 2$. Then, $\mathcal{S}_{\mathbb{C}}$ is generically r -identifiable if $r < \frac{\Pi}{\Sigma+1}$, unless (n_1, n_2, \dots, n_d) is*

⁴If one prefers to think projectively: two rank decompositions are essentially equal if the points on the Segre variety in projective space are equal. Note that we can ignore multiplicities of the points because a decomposition in which the same rank-1 tensor appears multiple times is not of rank r , contradicting the assumption on \mathfrak{A} 's rank.

⁵The letters are mnemonics for Diagonal scaling, Block diagonal scaling, the Permutation indeterminacies, and the Trivial indeterminacies.

⁶That is, \mathbf{q} is in the orbit of \mathbf{p} .

- (1) $(4, 4, 3)$,
- (2) $(4, 4, 4)$,
- (3) $(6, 6, 3)$,
- (4) $(n, n, 2, 2)$ with $n \in \mathbb{N}$,
- (5) $(2, 2, 2, 2, 2)$, or
- (6) $n_1 > \prod_{k=2}^d n_k - \sum_{k=2}^d (n_k - 1)$.

A computer-assisted proof [26] shows that the conjecture is true for all $n_1 \cdots n_d \leq 15000$. The conjecture is nearly optimal in the sense that a generic $\mathbf{p} \in \mathbb{C}^{r(\Sigma+d)}$ with $r > \frac{\Pi}{\Sigma+1}$ has a solution fiber at $\mathfrak{A} = f(\mathbf{p})$ of positive dimension, i.e., $\dim f^\dagger(\mathfrak{A}) \geq r(\Sigma+1) - \Pi > 0$, so that it is not identifiable. For “perfect” spaces where $r = \frac{\Pi}{\Sigma+1}$ and for which $\sigma_r(\mathcal{S}_{\mathbb{C}})$ is nondefective, a generic $\mathfrak{A} \in \sigma_r(\mathcal{S}_{\mathbb{C}}) = \mathbb{C}^\Pi$ has $|f^\dagger(\mathfrak{A})/\sim| = k$, where k is expected to be strictly larger than 1 according to [44].

Closely related with generic r -identifiability is the a priori weaker notion of *robust r -identifiability*, which was recently explored in [7].

Definition 1 (Robust r -identifiability). A rank- r tensor $\mathfrak{A} \in \sigma_r^0(\mathcal{S}_{\mathbb{F}})$ is said to be robustly r -identifiable if it is a regular point of $\sigma_r^0(\mathcal{S}_{\mathbb{F}})$ and there exists an open neighborhood \mathcal{N} of \mathfrak{A} such that

$$\forall \mathfrak{B} \in \mathcal{N} \subset \sigma_r^0(\mathcal{S}_{\mathbb{F}}) : |f^\dagger(\mathfrak{B})/\sim| = 1.$$

Notice that the assumption on the regularity of \mathfrak{A} ensures that Terracini’s matrix coincides with the Zariski tangent space and that (an open semi-algebraic subset of) the open neighborhood \mathcal{N} is Nash diffeomorphic with a semi-algebraic subset of $\mathbb{F}^{r(\Sigma+1)}$ [13, Section 2.9]. *Robust r -identifiability is a most convenient technical assumption for defining a condition number of the tensor rank decomposition problem. In fact, if a tensor is not robustly r -identifiable, then I will define the condition number to be ∞ .* Fortunately, the technical assumption on robust r -identifiability is very mild because of the following result.

Lemma 3. *Let $\mathcal{S}_{\mathbb{F}}$ be a generically r -identifiable Segre variety, and let Z be the Zariski-closed locus where r -identifiability fails. Then, every $\mathfrak{A} \in \sigma_r^0(\mathcal{S}_{\mathbb{F}}) \setminus Z$ is robustly r -identifiable.*

Proof. Since $\sigma_r(\mathcal{S}_{\mathbb{F}})$ lives in a Euclidean space there exists a minimum distance $\delta > 0$ from a fixed point $\mathfrak{A} \in \sigma_r^0(\mathcal{S}_{\mathbb{F}}) \setminus Z$ to the Euclidean-closed set Z . Taking the Euclidean-open ball \mathcal{B}_δ of radius δ around \mathfrak{A} by definition does not intersect Z . Hence, $\mathcal{B}_\delta \cap \sigma_r^0(\mathcal{S}_{\mathbb{F}}) \subset \sigma_r^0(\mathcal{S}_{\mathbb{F}}) \setminus Z$ contains only r -identifiable points and the radius $\delta > 0$ can be chosen so that it contains only one connected component, concluding the proof. \square

One may wonder why I choose to introduce this technical restriction for defining a condition number. After all, condition numbers could also be defined for set-valued maps g by defining a reasonable premetric for measuring distances between sets. I believe there are some reasons that justify the approach developed in this paper. The first reason is very pragmatic and common in (semi)algebraic geometry: the tensors that are not robustly r -identifiable are expected to be contained in a semi-algebraic set of codimension at least 1 in $\sigma_r(\mathcal{S}_{\mathbb{F}})$ by Conjecture 2 and Lemma 3, so that the definition is valid generically. In other words, if $\mathcal{S}_{\mathbb{F}}$ is a generically r -identifiable Segre variety and one samples a random rank- r decomposition $\mathfrak{A} \in \sigma_r^0(\mathcal{S}_{\mathbb{F}})$ from any probability distribution whose support is not contained in the Zariski-closed locus Z where r -identifiability fails, then \mathfrak{A} is both robustly r -identifiable and its condition number is well-defined with probability 1. Another reason is that the condition number admits a completely straightforward interpretation for robustly r -identifiable tensors: it asymptotically measures the ratio of the forward error and the backward error as in (3) with respect to some premetric that will be defined in the next section. It admits this straightforward interpretation because $f^\dagger(\mathfrak{A})/\sim$ locally behaves like an invertible function when \mathfrak{A} is robustly r -identifiable. A third motivation concerns the technique that I employed for defining the condition number in Theorem 7: it is inherently local. Hence, if \mathfrak{A} is robustly r -identifiable, a local analysis of one of the solutions $\mathbf{x} \in f^\dagger(\mathfrak{A})$ suffices for understanding the global behavior of all solutions $f^\dagger(\mathfrak{A})$ in a neighborhood of \mathfrak{A} . Finally, since many $\sigma_r(\mathcal{S}_{\mathbb{F}})$ are generically r -identifiable, it is sensible to redefine the tensor rank decomposition problem as the problem of finding the *unique* element in $f^\dagger(\mathfrak{A})/\sim$ for a given $\mathfrak{A} \in \sigma_r^0(\mathcal{S}_{\mathbb{F}})$. Then, the set of tensors

that admit multiple tensor rank decompositions are defined to be the *ill-posed* inputs. I should mention that this leads to a very natural alternative definition of a condition number, namely the minimum distance to an ill-posed input in the spirit of, e.g., [15, 34]. However, I will not pursue this interpretation of the proposed condition number in the present paper.

4. A SUITABLE DISTANCE MEASURE

In interpreting the N th largest singular value of Terracini's matrix as a condition number, as is stated in Theorem 1, understanding the influence of the part of the kernel given by (5) is crucial. Consider the following concrete situation. Let $b = (\mathbf{a}, \mathbf{b}, \mathbf{c}) \in \mathbb{F}^{n_1} \times \mathbb{F}^{n_2} \times \mathbb{F}^{n_3}$ be a representative of the rank-1 tensor $f(\mathbf{p}) = \mathbf{a} \otimes \mathbf{b} \otimes \mathbf{c}$, where $\mathbf{p} = \text{vec}(b)$. Suppose that the vectorized factor matrices \mathbf{p} are perturbed by the vector $\Delta_\epsilon = [\epsilon \mathbf{a}^T \quad -\epsilon \mathbf{b}^T \quad 0]^T$, which is contained in the span of K , where K is the specific basis in (5). Then,

$$\mathbf{p} = \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \\ \mathbf{c} \end{bmatrix} \quad \text{and} \quad \mathbf{p}_\epsilon = \mathbf{p} + \Delta_\epsilon = \begin{bmatrix} (1 + \epsilon)\mathbf{a} \\ (1 - \epsilon)\mathbf{b} \\ \mathbf{c} \end{bmatrix}.$$

For simplicity, let us assume that $\|\mathbf{a}\| = \|\mathbf{b}\| = \|\mathbf{c}\| = 1$. Then,

$$\|\mathbf{p} - \mathbf{p}_\epsilon\| = \sqrt{2}\epsilon \quad \text{while} \quad \|f(\mathbf{p}) - f(\mathbf{p}_\epsilon)\| = \|\mathbf{a} \otimes \mathbf{b} \otimes \mathbf{c} - (1 + \epsilon)(1 - \epsilon)\mathbf{a} \otimes \mathbf{b} \otimes \mathbf{c}\| = \epsilon^2.$$

This implies that if we would consider the condition number (2) with respect to the usual Euclidean norm, then we find that this particular perturbation bounds the relative condition number of the rank decomposition problem at \mathbf{p} from below by

$$\kappa_{2,2} \geq \lim_{\epsilon \rightarrow 0} \frac{\|\mathbf{p} - \mathbf{p}_\epsilon\| \cdot \|f(\mathbf{p})\|}{\|f(\mathbf{p}) - f(\mathbf{p}_\epsilon)\| \cdot \|\mathbf{p}\|} = \lim_{\epsilon \rightarrow 0} \frac{\sqrt{2}}{\sqrt{3}} \epsilon^{-1} \rightarrow \infty.$$

That is, a perturbation in the direction of the kernel results in an unbounded condition number, as was to be expected because the kernel originates from the group action of \mathcal{T} . Perturbations $\Delta_\epsilon \in \text{range}(K)$ of sufficiently small norm thus behave like a rescaling. Since rank-1 tensors are determined uniquely up to the action of \mathcal{T} , it seems a natural desire to ignore this action. Acting on this idea, suppose that we choose a different representative of $f(\mathbf{p}_\epsilon)$ by exploiting the additional freedom offered by \mathcal{T} . Consider the following representative

$$\mathbf{p}'_\epsilon = \text{diag}(1 - \epsilon, (1 - \epsilon)^{-1}, 1)\mathbf{p}_\epsilon = \begin{bmatrix} (1 - \epsilon^2)\mathbf{a} \\ \mathbf{b} \\ \mathbf{c} \end{bmatrix} \in \mathcal{T}\mathbf{p}_\epsilon.$$

Now $\|f(\mathbf{p}) - f(\mathbf{p}'_\epsilon)\| = \epsilon^2$, while $\|\mathbf{p} - \mathbf{p}'_\epsilon\| = \epsilon^2$, implying that the condition number should be

$$\kappa \geq \lim_{\epsilon \rightarrow 0} \frac{\epsilon^2}{\epsilon^2 \sqrt{3}} = \frac{1}{\sqrt{3}}.$$

It will be shown in Proposition 12 that the relative condition number with respect to the distance measure that will be introduced shortly is, in fact, equal to $\frac{1}{\sqrt{3}}$.

The foregoing example essentially illustrates that we cannot meaningfully define the condition number of the tensor rank decomposition *with respect to the standard Euclidean distance*. With this metric, the condition number is always unbounded. However, the example also suggests that exploiting the action of \mathcal{T} is crucial for obtaining a bounded condition number. I propose employing the following measure of distance for defining a condition number of the tensor rank decomposition.

Definition 2. Let $b_i = (\mathbf{a}_i^1, \mathbf{a}_i^2, \dots, \mathbf{a}_i^d)$ and $b'_i = (\mathbf{c}_i^1, \mathbf{c}_i^2, \dots, \mathbf{c}_i^d)$ be two set of representatives of the rank-1 tensors $p_i = \text{Seg}(b_i)$ and $q_i = \text{Seg}(b'_i)$ respectively for $i = 1, \dots, r$. Let $\mathbf{p} = \text{vecr}(b_1, b_2, \dots, b_r)$ and $\mathbf{q} = \text{vecr}(b'_1, b'_2, \dots, b'_r)$. The distance $d(\mathbf{p}, \mathbf{q})$ between the vectorized factor matrices \mathbf{p} and \mathbf{q} is defined as the least Euclidean distance between \mathbf{p} and the \mathcal{T} -orbit of \mathbf{q} :

$$d(\mathbf{p}, \mathbf{q}) = \inf_{T \in \mathcal{T}} \|\mathbf{p} - T\mathbf{q}\|.$$

Proposition 4. *Let b_i, b'_i, \mathbf{p} and \mathbf{q} be as in Definition 2. Then, the infimum is always attained:*

$$d(\mathbf{p}, \mathbf{q}) = \min_{T \in \mathcal{T}} \|\mathbf{p} - T\mathbf{q}\| \geq 0.$$

Moreover, the distance $d(\mathbf{p}, \mathbf{q}) = d(\mathbf{q}, \mathbf{p}) = 0$ if and only if $\mathbf{p} \sim \mathbf{q}$.

Proof. First note that $g(T) = \|\mathbf{p} - T\mathbf{q}\|$ is a coercive function in the parameters defining $T \in \mathcal{T}$, so that for a fixed input the minimizer of the optimization problem is always attained. Indeed, since $I \in \mathcal{T}$, it follows that the optimal $(P \otimes I_{\Sigma+d})B \in \mathcal{T}$ satisfies

$$\begin{aligned} \|\mathbf{p} - (P \otimes I_{\Sigma+d})B\mathbf{q}\|^2 &= \sum_{i=1}^r \|\text{vec}(b_i) - D_{\pi_i} \text{vec}(b'_{\pi_i})\|^2 \\ &= \sum_{i=1}^r \sum_{k=2}^d \|\mathbf{a}_i^k - \theta_{k,\pi_i} \mathbf{c}_{\pi_i}^k\|^2 + \sum_{i=1}^r \|\mathbf{a}_i^1 - (\theta_{2,\pi_i} \cdots \theta_{d,\pi_i})^{-1} \mathbf{c}_{\pi_i}^1\|^2 \\ &\leq \|\mathbf{p} - \mathbf{q}\|^2 < \infty \end{aligned}$$

where $B = \text{diag}(D_1, \dots, D_r)$, $D_i = \text{diag}(\theta_{2,i}^{-1} \cdots \theta_{d,i}^{-1} I_{n_1}, \theta_{2,i} I_{n_2}, \dots, \theta_{d,i} I_{n_d})$ and π is the permutation represented by the permutation matrix $P \in \mathbb{F}^{r \times r}$. Since we have a sum of positive reals, it follows that all $|\theta_{k,i}|$ are uniformly bounded from above by some positive constant C , and, hence there is also a uniform lower bound $0 < c$. So we may optimize $\theta_{k,i}$ over the bounded intervals $[-C, -c] \cup [c, C]$ leading to the conclusion that the infimum is attained.

The claim about $d(\mathbf{p}, \mathbf{q}) = 0$ is trivial. \square

Note in the foregoing proposition that if $\mathbf{q}' \in \mathcal{T}\mathbf{q}$ then $d(\mathbf{p}, \mathbf{q}') = d(\mathbf{p}, \mathbf{q})$. In particular, if \mathfrak{A} is an r -identifiable tensor, then the slightly abused notation $d(\mathbf{p}, f^\dagger(\mathfrak{A}))$ is well-defined.

The distance measure in Definition 2 may seem artificial, however it naturally measures the error between a given set of vectorized factor matrices $\mathbf{p} \in \mathbb{F}^{r(\Sigma+d)}$ and a solution of the rank decomposition problem—whose solution is only determined up to the action of \mathcal{T} . The `cpderr` function in the popular Tensorlab [68] software package for Matlab, for instance, computes the error in a similar way. Let $p_i = (\mathbf{a}_i^1, \mathbf{a}_i^2, \dots, \mathbf{a}_i^d)$ and $q_i = (\mathbf{b}_i^1, \mathbf{b}_i^2, \dots, \mathbf{b}_i^d)$, $i = 1, \dots, r$, be two set of representatives, and then let $A^k = [\mathbf{a}_j^k]_{j=1}^r$ and $B^k = [\mathbf{b}_j^k]_{j=1}^r$ be the corresponding set of factor matrices, where $k = 1, \dots, d$. Then, Tensorlab measures the error between two sets of factor matrices (A^1, \dots, A^d) and (B^1, \dots, B^d) as follows: it returns a d -tuple $(\xi_1, \xi_2, \dots, \xi_d)$ in which $\xi_i = \|A^i - B^i D_i P\|$, where D_i is a diagonal matrix and P is a permutation matrix. Tensorlab does not enforce $D_1 D_2 \dots D_d = I_r$, however if this additional constraint is imposed then a natural measures of the overall error, treating each of the factors equally, is $\sqrt{\xi_1^2 + \xi_2^2 + \dots + \xi_d^2} = d(\mathbf{a}, \mathbf{b})$, where $\mathbf{a} = \text{vecr}(p_1, p_2, \dots, p_r)$ and $\mathbf{b} = \text{vecr}(q_1, q_2, \dots, q_r)$.

Before concluding this section, it is natural to wonder whether we could eliminate the scaling indeterminacies on *both* representatives in Definition 2 by considering

$$\widehat{d}(\mathbf{p}, \mathbf{q}) = \inf_{T_1, T_2 \in \mathcal{T}} \|T_1 \mathbf{p} - T_2 \mathbf{q}\|,$$

which is essentially a premetric on the quotient space $\mathbb{F}^{r(\Sigma+1)}/\mathcal{T}$ and has some additional appeal because at least it is symmetric, i.e., $\widehat{d}(\mathbf{p}, \mathbf{q}) = \widehat{d}(\mathbf{q}, \mathbf{p})$. However, I believe that this particular choice will not lead to a sensible definition of a condition number because the Hausdorff separation property is lost, i.e., $\widehat{d}(\mathbf{p}, \mathbf{q}) = 0$ even when $\mathbf{p} \not\sim \mathbf{q}$ and $f(\mathbf{p}) \neq f(\mathbf{q})$. The reason is that the set over which we optimize is not closed, and this time the two matrices can counteract each other. The simplest example of this phenomenon are the following representatives of two distinct rank-1 tensors

$$\mathbf{p} = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \mathbf{c} \end{bmatrix} \quad \text{and} \quad \mathbf{q} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \mathbf{c} \end{bmatrix}.$$

Choosing $T_1 = T_2 = \text{diag}(\theta, \theta, \theta^{-2})$, we find

$$\lim_{\theta \rightarrow 0} \left\| \begin{bmatrix} \theta \mathbf{a}_1 - \theta \mathbf{b}_1 \\ \theta \mathbf{a}_2 - \theta \mathbf{b}_2 \\ \theta^{-2} \mathbf{c} - \theta^{-2} \mathbf{c} \end{bmatrix} \right\| = \lim_{\theta \rightarrow 0} \left\| \begin{bmatrix} \theta \mathbf{a}_1 - \theta \mathbf{b}_1 \\ \theta \mathbf{a}_2 - \theta \mathbf{b}_2 \\ 0 \end{bmatrix} \right\| \rightarrow 0,$$

hence $\widehat{d}(\mathbf{p}, \mathbf{q}) = 0$ while $f(\mathbf{p}) = \mathbf{a}_1 \otimes \mathbf{a}_2 \otimes \mathbf{c} \neq \mathbf{b}_1 \otimes \mathbf{b}_2 \otimes \mathbf{c} = f(\mathbf{q})$. That is, the orbits $\mathcal{T}\mathbf{p}$ and $\mathcal{T}\mathbf{q}$ meet at infinity.

5. THE ITERATED SCALING LEMMA

We are now in a position to state the main technical result in connection to Terracini's matrix that we will need for deriving the condition number with respect to the distance measure in Definition 2. It states that a perturbation of a representative p_i of a rank-1 tensor in the direction of the kernel of Terracini's matrix may always be interpreted as choosing a new representative of that rank-1 tensor $\text{Seg}(p_i) \in \mathcal{S}_{\mathbb{F}}$ plus a small perturbation Δ that is contained in the column span of Terracini's matrix. In the formulation and proof of the following lemma, I will slightly abuse notation for brevity, writing $q \in \mathcal{D}p$ when I mean $\text{vec}(q) \in \mathcal{D} \text{vec}(p)$ with $q, p \in \mathbb{F}^{n_1} \times \dots \times \mathbb{F}^{n_d}$.

Lemma 5 (Iterated Scaling). *Let $\mathcal{S}_{\mathbb{F}}$ be a Segre variety. For $i = 1, 2, \dots, r$, let*

$$p_i = (\mathbf{a}_i^1, \dots, \mathbf{a}_i^d), \quad \nabla_i = (\mathbf{n}_i^1, \dots, \mathbf{n}_i^d), \quad \text{and} \quad q_i = p_i + \nabla_i = (\mathbf{a}_i^1 + \mathbf{n}_i^1, \dots, \mathbf{a}_i^d + \mathbf{n}_i^d),$$

where $\text{Seg}(p_i), \text{Seg}(q_i) \in \mathcal{S}_{\mathbb{F}}$. Assume that the perturbation $\nabla = \text{vecr}(\nabla_1, \nabla_2, \dots, \nabla_r)$ is of sufficiently small norm:

$$(7) \quad \|\nabla\| \leq \frac{1}{2} \lambda^{-1} = \frac{1}{2^{d+4}(d-1)^{3/2}} \cdot \min_{1 \leq i \leq r} \min_{2 \leq k \leq d} (\|\mathbf{a}_i^1\|^{-1} \|\mathbf{a}_i^k\|^2).$$

If Terracini's matrix $T_{\mathbf{p}}$ associated with the vectorized factor matrices $\mathbf{p} = \text{vecr}(p_1, \dots, p_r)$ has rank $r(\Sigma + 1)$ and if ∇ is contained in its kernel, i.e.,

$$T_{\mathbf{p}} \nabla = 0,$$

then there is a representative $\dot{p}_i \in \mathcal{D}p_i$ such that $p_i + \nabla_i = q_i = \dot{p}_i + \Delta_i$ where $\Delta_i = \text{vec}(\Delta_i) \in \text{span}(K_i)^\perp$, and $\|\Delta_i\| \leq 2\lambda \|\nabla_i\|^2 \leq \|\nabla_i\|$ with $\nabla_i = \text{vec}(\nabla_i)$ and K_i as in (5). In other words, there exists a factorization

$$\mathbf{p} + \nabla = \dot{\mathbf{p}} + \Delta,$$

for which

$$\dot{\mathbf{p}} \sim \mathbf{p}, \quad \Delta \in \text{span}(K)^\perp, \quad \text{and} \quad \|\Delta\| \leq 2\lambda \|\nabla\|^2 \leq \|\nabla\|,$$

where K is as in (5).

Proof. We prove the assertion by proving the existence of a linearly convergent sequence of representatives $p_i^{(k)} \in \mathcal{D}p_i$, $k = 1, 2, \dots$, with the following properties

$$(H0a) \quad q_i - p_i^{(k)} = \Delta_i^{(k)} + \nabla_i^{(k)},$$

where

$$\Delta_i^{(k)} = \text{vec}(\Delta_i^{(k)}) \in \text{span}(K_i)^\perp \quad \text{and} \quad \nabla_i^{(k)} = \text{vec}(\nabla_i^{(k)}) \in \text{span}(K_i).$$

The sequence will be constructed in such a way that

$$(H0b) \quad \lim_{k \rightarrow \infty} \|\nabla_i^{(k)}\| \rightarrow 0.$$

As $\nabla_i^{(k)}$ is contained in a finite-dimensional subspace of $\mathbb{F}^{\Sigma+d}$, it follows immediately that $\nabla_i^{(k)} \rightarrow 0$ as $k \rightarrow \infty$. Under the assumptions of the lemma both $\mathbf{p}_i^{(k)} = \text{vec}(p_i^{(k)})$ and $\Delta_i^{(k)}$ will be shown to be of uniformly bounded norm, so that at least a convergent subsequence exists for which both

$$\lim_{k \rightarrow \infty} \Delta_i^{(k)} \rightarrow \Delta_i = \text{vec}(\Delta_i) \quad \text{and} \quad \lim_{k \rightarrow \infty} \mathbf{p}_i^{(k)} \rightarrow \dot{\mathbf{p}}_i = \text{vec}(\dot{p}_i)$$

are well defined. So it remains to prove that a sequence satisfying hypotheses (H0a) and (H0b) exists.

Part I: A recurrence relation. Let $p_i^{(1)} = p_i$ and $\Delta_i^{(1)} = 0$, and define $\nabla_i^{(1)} = \nabla_i$. Since $\nabla^{(1)} = \text{vecr}(\nabla_1^{(1)}, \dots, \nabla_r^{(1)}) = \nabla$ is contained in the kernel, we have

$$T_{\mathbf{p}} \nabla^{(1)} = \sum_{i=1}^r T_i \nabla_i^{(1)} = 0,$$

where $T_{\mathbf{p}}$ and T_i are as in (4). Suppose that $T_i \nabla_i^{(1)} \neq 0$, so that $\nabla_i^{(1)}$ would not be contained in the span of K_i . Then $T_{\mathbf{p}}$'s rank would be strictly less than the expected value $r(\Sigma + 1)$, which is a contradiction. Consequently, the base case $k = 1$ of (H0a) is true. Assume now that the statement is true for all $l = 1, 2, \dots, k$, and then we show that it holds for $k + 1$ as well. Since K_i is a basis of the kernel of T_i , we can express

$$(8) \quad \nabla_i^{(k)} = [\mathbf{k}_i^2 \quad \dots \quad \mathbf{k}_i^d] \mathbf{v}_i^{(k)} = K_i \mathbf{v}_i^{(k)}$$

for some $\mathbf{v}_i^{(k)} \in \mathbb{F}^{d-1}$. By hypothesis (H0a), $p_i^{(k)} \in \mathcal{D}p_i$, so that we can write it explicitly as

$$(9) \quad p_i^{(k)} = (\gamma_{1,i}^{(k)} \mathbf{a}_i^1, \gamma_{2,i}^{(k)} \mathbf{a}_i^2, \dots, \gamma_{d,i}^{(k)} \mathbf{a}_i^d);$$

clearly this is true with $\gamma_{j,i}^{(1)} = 1$ for the base case, and will follow shortly for the induction step as well. Consider the perturbed representative $p_i^{(k)} + \nabla_i^{(k)}$. By writing this perturbation $\nabla_i^{(k)}$ with respect to the particular basis K_i of the kernel of T_i as in (8), we find that

$$(10) \quad z_i^{(k)} = p_i^{(k)} + \nabla_i^{(k)} = \left((\gamma_{1,i}^{(k)} + v_{j,i}^{(k)} + \dots + v_{d,i}^{(k)}) \mathbf{a}_i^1, (\gamma_{2,i}^{(k)} - v_{2,i}^{(k)}) \mathbf{a}_i^2, \dots, (\gamma_{d,i}^{(k)} - v_{d,i}^{(k)}) \mathbf{a}_i^d \right),$$

where $v_{j,i}^{(k)}$ denotes the $(j - 1)$ th element of $\mathbf{v}_i^{(k)}$. Consequently, $\text{Seg}(z_i^{(k)})$ and $\text{Seg}(p_i^{(k)})$ are linearly dependent, i.e., multiples of each other. I claim that the following representative

$$(11) \quad p_i^{(k+1)} = \left(\prod_{j=2}^d (\gamma_{j,i}^{(k)} - v_{j,i}^{(k)})^{-1} \mathbf{a}_i^1, (\gamma_{2,i}^{(k)} - v_{2,i}^{(k)}) \mathbf{a}_i^2, \dots, (\gamma_{d,i}^{(k)} - v_{d,i}^{(k)}) \mathbf{a}_i^d \right) \in \mathcal{D}p_i$$

will induce the required sequence. Define

$$(12) \quad z_i^{(k)} - p_i^{(k+1)} = \left(\left(\gamma_{1,i}^{(k)} + \sum_{j=2}^d v_{j,i}^{(k)} - \prod_{j=2}^d (\gamma_{j,i}^{(k)} - v_{j,i}^{(k)})^{-1} \right) \mathbf{a}_i^1, 0, \dots, 0 \right) = \nabla_i^{(k+1)} + \widehat{\Delta}_i^{(k+1)},$$

where the factorization is such that $\nabla_i^{(k+1)} \in \text{span}(K_i)$ and $\text{vec}(\widehat{\Delta}_i^{(k+1)}) \in \text{span}(K_i)^\perp$. Then,

$$z_i^{(k)} = p_i^{(k)} + \nabla_i^{(k)} = p_i^{(k+1)} + \nabla_i^{(k+1)} + \widehat{\Delta}_i^{(k+1)}.$$

Since this must be inductively true for $l = 1, 2, \dots, k$, we find that

$$(13) \quad q_i = z_i^{(1)} = p_i^{(k+1)} + \nabla_i^{(k+1)} + \sum_{j=2}^{k+1} \widehat{\Delta}_i^{(j)} = p_i^{(k+1)} + \nabla_i^{(k+1)} + \Delta_i^{(k+1)}.$$

This proves (H0a).

Part II: Convergence. It only remains to show (H0b), i.e., the sequence of $\nabla_i^{(k)}$'s converges to the zero vector. We demonstrate that

$$(H1) \quad \|\nabla_i^{(k)}\| \leq \lambda^{k-1} \|\nabla_i\|^k \quad \text{for some constant } \lambda \in \mathbb{R}, 0 < \lambda;$$

for the base case $k = 1$ this is readily true as $\nabla_i^{(1)} = \nabla_i$ by definition. We may assume the truth of the above statement for $l = 1, 2, \dots, k$. From (12) it immediately follows that

$$\|\nabla_i^{(k+1)}\| \leq \|\mathbf{z}_i^{(k)} - \mathbf{p}_i^{(k+1)}\| = \|\mathbf{a}_i^1\| \cdot \left| \gamma_{1,i}^{(k)} + \sum_{j=2}^d v_{j,i}^{(k)} - \prod_{j=2}^d (\gamma_{j,i}^{(k)} - v_{j,i}^{(k)})^{-1} \right|,$$

where $\mathbf{z}_i^{(k)} = \text{vec}(z_i^{(k)})$. Our efforts will be focused on showing that the right hand side is bounded by $\lambda^k \|\nabla_i\|^{k+1}$. As an intermezzo, note that the boundedness of $\|\Delta_i^{(k+1)}\|$ would then follow from the observation that

$$\|\Delta_i^{(k+1)}\| \leq \sum_{j=2}^{k+1} \|\widehat{\Delta}_i^{(j)}\| \leq \sum_{j=1}^k \|\mathbf{z}_i^{(j)} - \mathbf{p}_i^{(j+1)}\| \leq \sum_{j=1}^k \lambda^j \|\nabla_i\|^{j+1} \leq \lambda^{-1} \sum_{j=1}^{\infty} (\lambda \|\nabla_i\|)^{j+1} = \lambda \frac{\|\nabla_i\|^2}{1 - \lambda \|\nabla_i\|},$$

provided that $\lambda \|\nabla_i\| < 1$ but this will be satisfied by our assumptions. The boundedness of $\mathbf{p}_i^{(k)}$ is then an immediate consequence of (13) and the uniform boundedness of $\mathbf{q}_i = \text{vec}(q_i)$, $\nabla_i^{(k+1)}$ and $\Delta_i^{(k+1)}$. In fact, since we will assume (B3), i.e., $\lambda \|\nabla_i\| \leq \frac{1}{2}$, it follows that $\|\Delta_i^{(k+1)}\| \leq 2 \cdot \frac{1}{2} \|\nabla_i\|$, which already proves the bound on $\|\Delta_i\|$.

For bounding $\|\mathbf{z}_i^{(k)} - \mathbf{p}_i^{(k+1)}\|$ from above, we proceed as follows. During our derivations, we will assume some additional convenient constraints on certain quantities; they will be considered more carefully in the next part. Consider the coefficient of \mathbf{a}_i^1 in (12). First, a bound on the $v_{j,i}^{(k)}$'s is obtained as follows. Since $K_i \in \mathbb{F}^{(\Sigma+d) \times (d-1)}$ is a basis for the kernel of T_i it has linearly independent columns. Therefore, the Moore–Penrose pseudoinverse is $K_i^\dagger = (K_i^H K_i)^{-1} K_i^H$, and as $\nabla_i^{(k)} \in \text{span}(K_i)$ one has $\nabla_i^{(k)} = K_i K_i^\dagger \nabla_i^{(k)}$, because $K_i K_i^\dagger$ is a projector onto the column span of K_i ; so $\mathbf{v}_i^{(k)} = K_i^\dagger \nabla_i^{(k)}$. If $\varsigma_i(A)$ denotes the i th largest singular value of $A \in \mathbb{F}^{m \times n}$ and $\lambda_i(A)$ denotes the i th largest eigenvalue of a Hermitian matrix $A \in \mathbb{F}^{m \times m}$, then it is well-known that

$$\|K_i^\dagger\|_2 = \varsigma_1(K_i^\dagger) = (\varsigma_{d-1}(K_i))^{-1} = (\lambda_{d-1}(K_i^H K_i))^{-\frac{1}{2}}.$$

One may verify by direct computation that

$$K_i^H K_i = \text{diag}(\|\mathbf{a}_i^2\|^2, \|\mathbf{a}_i^3\|^2, \dots, \|\mathbf{a}_i^d\|^2) + \|\mathbf{a}_i^1\|^2 \mathbf{1} \mathbf{1}^T,$$

where $\mathbf{1}$ is a vector of length $d-1$ containing only ones. From [81, Section 2.41] it follows that the eigenvalues of $K_i^H K_i$ satisfy

$$\lambda_j(K_i^H K_i) \geq \lambda_j(\text{diag}(\|\mathbf{a}_i^2\|^2, \|\mathbf{a}_i^3\|^2, \dots, \|\mathbf{a}_i^d\|^2)), \quad j = 1, 2, \dots, d-1,$$

so in particular we can conclude that

$$(14) \quad \|K_i^\dagger\|_2 \leq \left(\min_{2 \leq k \leq d} \|\mathbf{a}_i^k\|^2 \right)^{-\frac{1}{2}} = \chi_i.$$

We can now bound the $(j-1)$ th element $v_{j,i}^{(k)}$, $j = 2, 3, \dots, d$, of $\mathbf{v}_i^{(k)}$ as follows:

$$(15) \quad |v_{j,i}^{(k)}| \leq \|\mathbf{v}_i^{(k)}\| \leq \|K_i^\dagger\|_2 \|\nabla_i^{(k)}\| \leq \chi_i \|\nabla_i^{(k)}\| \leq \chi_i \lambda^{k-1} \|\nabla_i\|^k,$$

where the last step is due to the induction hypothesis (H1).

For convenience, we define

$$\epsilon_{j,i}^{(k)} = \sum_{\ell=1}^k v_{j,i}^{(\ell)} \quad \text{and} \quad \epsilon_{j,i}^{(0)} = 0,$$

so that

$$(16) \quad |\epsilon_{j,i}^{(k)}| \leq \sum_{\ell=1}^k |v_{j,i}^{(\ell)}| \leq \chi_i \lambda^{-1} \sum_{\ell=1}^k (\lambda \|\nabla_i\|)^\ell \leq \chi_i \lambda^{-1} \sum_{\ell=1}^{\infty} (\lambda \|\nabla_i\|)^\ell = \frac{\chi_i \|\nabla_i\|}{1 - \lambda \|\nabla_i\|} = C'_i;$$

the second inequality is by the induction hypothesis, and the penultimate equality holds for $\lambda \|\nabla_i\| < 1$, which will be satisfied by our assumptions. We additionally assume the following bound

$$(B0) \quad C'_i \leq \frac{1}{2(d-1)},$$

which will be investigated in more detail later. Then, for $j = 2, 3, \dots, d$ we find from (9) and (11) that

$$(17a) \quad \gamma_{j,i}^{(k+1)} = \gamma_{j,i}^{(k)} - v_{j,i}^{(k)} = \dots = \gamma_{j,i}^{(1)} - \sum_{\ell=1}^k v_{j,i}^{(\ell)} = 1 - \epsilon_{j,i}^{(k)}, \quad j = 2, 3, \dots, d.$$

By induction, the above formula also applies for $\gamma_{j,i}^{(\ell)}$ with $\ell = 1, 2, \dots, k$. For $j = 1$ we have

$$(17b) \quad \gamma_{1,i}^{(k+1)} = \prod_{j=2}^d (\gamma_{j,i}^{(k)} - v_{j,i}^{(k)})^{-1} = \prod_{j=2}^d (1 - \epsilon_{j,i}^{(k-1)} - v_{j,i}^{(k)})^{-1} = \prod_{j=2}^d (1 - \epsilon_{j,i}^{(k)})^{-1},$$

which by induction is also true for $\gamma_{1,i}^{(\ell)}$ with $\ell = 1, 2, \dots, k$. With the above observations, the coefficient of \mathbf{a}_i^1 in (12) may be written as

$$(18) \quad \gamma_{1,i}^{(k)} + \sum_{j=2}^d v_{j,i}^{(k)} - \gamma_{1,i}^{(k+1)} = \prod_{j=2}^d (1 - \epsilon_{j,i}^{(k-1)})^{-1} + \sum_{j=2}^d v_{j,i}^{(k)} - \prod_{j=2}^d (1 - \epsilon_{j,i}^{(k-1)} - v_{j,i}^{(k)})^{-1}.$$

For proceeding, we will need some fairly convoluted series expansions; however, the key idea revolves around expanding $(1-x)^{-1}$ with $x \approx 0$ by a Maclaurin series. Since we will be confronted with products of such expansions, it is worthwhile to recall the following general formula:

$$(19) \quad \prod_{j=1}^n (1 - x_j)^{-1} = \prod_{j=1}^n \sum_{\kappa=0}^{\infty} x_j^{\kappa} = \sum_{\kappa=0}^{\infty} \sum_{\|\ell\|_1=\kappa} \prod_{j=1}^n x_j^{\ell_j},$$

where $\ell = [\ell_1 \ \ell_2 \ \dots \ \ell_n] \in \mathbb{N}^n$ and $\|\ell\|_1 = \sum_{j=1}^n |\ell_j|$ is the 1-norm of ℓ ; in the expression it is assumed that $|x_j| < 1$ so that all expansions, including their products, are absolutely convergent. The last term in (18) can be rewritten as follows:

$$\begin{aligned} Z &= \prod_{j=2}^d (1 - \epsilon_{j,i}^{(k-1)} - v_{j,i}^{(k)})^{-1} = \prod_{\ell=2}^d (1 - \epsilon_{\ell,i}^{(k-1)})^{-1} \prod_{j=2}^d \frac{1 - \epsilon_{j,i}^{(k-1)}}{1 - \epsilon_{j,i}^{(k-1)} - v_{j,i}^{(k)}} \\ &= \prod_{\ell=2}^d (1 - \epsilon_{\ell,i}^{(k-1)})^{-1} \prod_{j=2}^d \left(1 - \frac{v_{j,i}^{(k)}}{1 - \epsilon_{j,i}^{(k-1)}}\right)^{-1} \\ &= \prod_{\ell=2}^d (1 - \epsilon_{\ell,i}^{(k-1)})^{-1} \sum_{\kappa=0}^{\infty} \sum_{\|\ell\|_1=\kappa} \prod_{j=2}^d \left(\frac{v_{j,i}^{(k)}}{1 - \epsilon_{j,i}^{(k-1)}}\right)^{\ell_j}, \end{aligned}$$

where the last step was by (19), which requires that

$$(B1) \quad |v_{j,i}^{(k)}| < |1 - \epsilon_{j,i}^{(k-1)}|;$$

this hypothesis will be investigated later. Let us write $Z = S_0 + S_1 + S_{\infty}$, where

$$\begin{aligned} S_0 &= \prod_{j=2}^d (1 - \epsilon_{j,i}^{(k-1)})^{-1} = \gamma_{1,i}^{(k)}, \\ S_1 &= \prod_{\ell=2}^d (1 - \epsilon_{\ell,i}^{(k-1)})^{-1} \cdot \sum_{j=2}^d \frac{v_{j,i}^{(k)}}{1 - \epsilon_{j,i}^{(k-1)}}, \text{ and} \\ S_{\infty} &= \prod_{\ell=2}^d (1 - \epsilon_{\ell,i}^{(k-1)})^{-1} \cdot \sum_{\kappa=2}^{\infty} \sum_{\|\ell\|_1=\kappa} \prod_{j=2}^d \left(\frac{v_{j,i}^{(k)}}{1 - \epsilon_{j,i}^{(k-1)}}\right)^{\ell_j}. \end{aligned}$$

We will need to apply (19) again to S_1 for obtaining the desired result.

$$\begin{aligned} S_1 &= \sum_{j=2}^d v_{j,i}^{(k)} \left(1 + \frac{\epsilon_{j,i}^{(k-1)}}{1 - \epsilon_{j,i}^{(k-1)}}\right) \cdot \prod_{\ell=2}^d (1 - \epsilon_{\ell,i}^{(k-1)})^{-1} \\ &= \sum_{j=2}^d v_{j,i}^{(k)} \left(1 + \frac{\epsilon_{j,i}^{(k-1)}}{1 - \epsilon_{j,i}^{(k-1)}}\right) \cdot \left(1 + \sum_{\kappa=1}^{\infty} \sum_{\|\ell\|_1=\kappa} \prod_{l=2}^d (\epsilon_{l,i}^{(k-1)})^{\ell_l}\right) \end{aligned}$$

$$= \underbrace{\sum_{j=2}^d v_{j,i}^{(k)}}_{S'_1} + \underbrace{\sum_{j=2}^d v_{j,i}^{(k)} \left(\sum_{\kappa=1}^{\infty} \sum_{\|\ell\|_1=\kappa} \prod_{l=2}^d (\epsilon_{l,i}^{(k-1)})^{\ell_l} \right)}_{S''_1} + \underbrace{\sum_{j=2}^d v_{j,i}^{(k)} \frac{\epsilon_{j,i}^{(k-1)}}{1 - \epsilon_{j,i}^{(k-1)}} \cdot \prod_{\ell=2}^d (1 - \epsilon_{\ell,i}^{(k-1)})^{-1}}_{S'''_1}.$$

It is now clear that

$$(18) = S_0 + S'_1 - Z = S_0 + S'_1 - (S_0 + S'_1 + S''_1 + S'''_1 + S_{\infty}) = -(S''_1 + S'''_1 + S_{\infty}).$$

If we could show that the last inequality in

$$|(18)| = |S''_1 + S'''_1 + S_{\infty}| \leq |S''_1| + |S'''_1| + |S_{\infty}| \leq \|\mathbf{a}_i^1\|^{-1} \lambda^k \|\nabla_i\|^{k+1}$$

holds, then the proof of (H1) would be concluded. We proceed along this line. From (16),

$$1 - \epsilon_{j,i}^{(k-1)} \geq 1 - |\epsilon_{j,i}^{(k-1)}| \geq 1 - C'_i,$$

and as $|\epsilon_{j,i}^{(k-1)}| \leq \frac{1}{2}$ by (B0), it follows that

$$|1 - \epsilon_{j,i}^{(k-1)}|^{-1} = (1 - \epsilon_{j,i}^{(k-1)})^{-1} \leq (1 - C'_i)^{-1} = C_i;$$

this bound is uniform in the sense that it depends neither on j nor k . Consequently,

$$\begin{aligned} \left| \sum_{\|\ell\|_1=\kappa} \prod_{j=2}^d \left(\frac{v_{j,i}^{(k)}}{1 - \epsilon_{j,i}^{(k-1)}} \right)^{\ell_j} \right| &\leq C_i^{\kappa} \sum_{\|\ell\|_1=\kappa} \prod_{j=2}^d |v_{j,i}^{(k)}|^{\ell_j} \leq C_i^{\kappa} \|\mathbf{v}_i^{(k)} \otimes \cdots \otimes \mathbf{v}_i^{(k)}\|_1 \\ (20) \qquad \qquad \qquad &\leq C_i^{\kappa} (d-1)^{\kappa/2} \|\mathbf{v}_i^{(k)}\|^{\kappa}, \end{aligned}$$

where the second step is because every ℓ with $\|\ell\|_1 = \kappa$ can be identified with the multi-index

$$(\underbrace{1, \dots, 1}_{\ell_1}, \underbrace{2, \dots, 2}_{\ell_2}, \dots, \underbrace{d-1, \dots, d-1}_{\ell_{d-1}})$$

whose length is κ . From this identification, it follows that the symmetric part of $\mathbf{v}_i^{(k)} \otimes \cdots \otimes \mathbf{v}_i^{(k)}$ contains precisely all summands of $\sum_{\|\ell\|_1=\kappa} \prod_{j=2}^d |v_{j,i}^{(k)}|^{\ell_j}$. As a result, S_{∞} can be bounded as follows

$$|S_{\infty}| \leq C_i^{d-1} \sum_{\kappa=2}^{\infty} (C_i \sqrt{d-1} \|\mathbf{v}_i^{(k)}\|)^{\kappa} \leq \frac{C_i^{d+1} (d-1) \chi_i^2 \lambda^{2k-2} \|\nabla_i\|^{2k}}{1 - C_i \sqrt{d-1} \chi_i \lambda^{k-1} \|\nabla_i\|^k},$$

where (15) was used to bound $\|\mathbf{v}_i^{(k)}\|$ in the second inequality; herein, we assumed that

$$(B2) \qquad C_i \sqrt{d-1} \chi_i \lambda^{k-1} \|\nabla_i\|^k \leq \frac{1}{2}$$

so that the series is convergent. In an analogous fashion as in the derivation of (20), one finds

$$(21) \quad \left| \sum_{\|\ell\|_1=\kappa} \prod_{j=2}^d (\epsilon_{j,i}^{(k-1)})^{\ell_j} \right| \leq \|\epsilon_i^{(k-1)} \otimes \cdots \otimes \epsilon_i^{(k-1)}\|_1 \leq (d-1)^{\kappa/2} \|\epsilon_i^{(k-1)}\|^{\kappa} \leq ((d-1)C'_i)^{\kappa},$$

where $\epsilon_i^{(k)} = [\epsilon_{2,i}^{(k-1)} \cdots \epsilon_{d,i}^{(k-1)}]$ and the last step is due to (16). Then, the bound for S''_1 is

$$|S''_1| \leq \sum_{j=2}^d |v_{j,i}^{(k)}| \cdot \left(\sum_{\kappa=1}^{\infty} ((d-1)C'_i)^{\kappa} \right) = \frac{(d-1)C'_i}{1 - (d-1)C'_i} \sum_{j=2}^d |v_{j,i}^{(k)}| = \frac{(d-1)C'_i}{1 - (d-1)C'_i} \|\mathbf{v}_i^{(k)}\|_1;$$

the Maclaurin series is convergent because of (B0). We continue with a bound for S'''_1 :

$$|S'''_1| \leq \sum_{j=2}^d |v_{j,i}^{(k)}| C'_i C_i \prod_{\ell=2}^d |1 - \epsilon_{\ell,i,k-1}|^{-1} \leq C'_i C_i^d \sum_{j=2}^d |v_{j,i}^{(k)}| = C'_i C_i^d \|\mathbf{v}_i^{(k)}\|_1.$$

It follows from the foregoing two bounds that

$$|S''_1| + |S'''_1| \leq \left(\frac{(d-1)C'_i}{1 - (d-1)C'_i} + C'_i C_i^d \right) \|\mathbf{v}_i^{(k)}\|_1 \leq C'_i \sqrt{d-1} \left(\frac{d-1}{1 - (d-1)C'_i} + C_i^d \right) \chi_i \lambda^{k-1} \|\nabla_i\|^k$$

$$\leq \left(\frac{d-1}{1-(d-1)C'_i} + C_i^d \right) \frac{\chi_i^2 \sqrt{d-1}}{1-\lambda \|\nabla_i\|} \lambda^{k-1} \|\nabla_i\|^{k+1},$$

where we used $\|\mathbf{v}_i^{(k)}\|_1 \leq \sqrt{d-1} \|\mathbf{v}_i^{(k)}\|$ and (15). Let $\zeta = \frac{d-1}{1-(d-1)C'_i} + C_i^d$. Combining the above with the bound for S_∞ , we get

$$|(18)| \leq \chi_i^2 \sqrt{d-1} \left(\frac{\zeta}{1-\lambda \|\nabla_i\|} + \frac{C_i^{d+1} \sqrt{d-1} \lambda^{k-1} \|\nabla_i\|^{k-1}}{1-C_i \sqrt{d-1} \chi_i \lambda^{k-1} \|\nabla_i\|^k} \right) \lambda^{k-1} \|\nabla_i\|^{k+1}.$$

Therefore, it suffices to demonstrate that the right hand side in the above expression is less than $\|\mathbf{a}_i^1\|^{-1} \lambda^k \|\nabla_i\|^{k+1}$. So, it suffices showing that

$$(22) \quad \|\mathbf{a}_i^1\| \chi_i^2 \sqrt{d-1} \left(\frac{\zeta}{1-\lambda \|\nabla_i\|} + \frac{C_i^{d+1} \sqrt{d-1} \lambda^{k-1} \|\nabla_i\|^{k-1}}{1-C_i \sqrt{d-1} \chi_i \lambda^{k-1} \|\nabla_i\|^k} \right) \leq \lambda.$$

Let us assume additionally that

$$(B3) \quad \lambda \|\nabla_i\| \leq \frac{1}{2}.$$

Then, by exploiting (B0), (B2), and (B3), it follows that (22) is implied by the following inequality

$$(B4) \quad \|\mathbf{a}_i^1\| \chi_i^2 \sqrt{d-1} (4(d-1) + 2^{d+1} + 2^{d-k+3} \sqrt{d-1}) \leq \lambda.$$

Note that $\lambda > 0$ is a free parameter, i.e., it is not specified as input to the lemma, so it can simply be chosen so as to satisfy the above inequality. For instance, let us choose

$$\lambda = 2^{d+3} (d-1)^{3/2} \cdot \max_{1 \leq i \leq r} (\chi_i^2 \|\mathbf{a}_i^1\|).$$

Provided that (B0), (B1), (B2), (B3) and (B4) are true, this proves (H1), which implies (H0b).

Part III: Eliminating assumptions. For concluding the proof, we show that (7) implies (B0), (B1), (B2) and (B3). The assumption (B0) can be eliminated as follows:

$$(B0) \quad \Leftrightarrow \quad \frac{\|\nabla_i\|}{1-\lambda \|\nabla_i\|} \leq \frac{1}{2(d-1)} \chi_i^{-1} \quad \Leftrightarrow \quad \left(\|\nabla_i\| \leq \frac{1}{4(d-1)} \chi_i^{-1} \right) \wedge (B3),$$

where \wedge denotes the logical conjunction. The last statement also implies (B1), because we have

$$\begin{aligned} (B1) \quad \Leftrightarrow \quad \chi_i \lambda^{k-1} \|\nabla_i\|^k &< 1 - \frac{\chi_i \|\nabla_i\|}{1-\lambda \|\nabla_i\|} \quad \Leftrightarrow \quad \left(\chi_i \|\nabla_i\| < 1 - \frac{\chi_i \|\nabla_i\|}{1-\lambda \|\nabla_i\|} \right) \wedge (B3) \\ &\Leftrightarrow \quad \left(\chi_i \|\nabla_i\| \left(1 + \frac{1}{1-\lambda \|\nabla_i\|} \right) < 1 \right) \wedge (B3) \\ &\Leftrightarrow \quad \left(\|\nabla_i\| \leq \frac{1}{3} \chi_i^{-1} \right) \wedge (B3). \end{aligned}$$

The elimination of (B2) proceeds as follows:

$$\begin{aligned} (B2) \quad \Leftrightarrow \quad \lambda^{k-1} \|\nabla_i\|^k &\leq \frac{1}{2\sqrt{d-1}C_i} \chi_i^{-1} \quad \Leftrightarrow \quad \left(\|\nabla_i\| \leq \frac{1}{4\sqrt{d-1}} \chi_i^{-1} \right) \wedge (B0) \wedge (B3) \\ &\Leftrightarrow \quad \left(\|\nabla_i\| \leq \frac{1}{4(d-1)} \chi_i^{-1} \right) \wedge (B3). \end{aligned}$$

It follows that (B3), or, equivalently, $\|\nabla_i\| \leq \frac{1}{2} \lambda^{-1}$ is the strongest bound, because we have

$$\|\nabla_i\| \leq \frac{1}{2} \lambda^{-1} = \frac{1}{2^{d+2} \sqrt{d-1}} \cdot \frac{1}{4(d-1)} \min_{1 \leq i \leq r} (\|\mathbf{a}_i^1\|^{-1} \chi_i^{-2}) < \frac{1}{4(d-1)} \chi_i^{-1} < \frac{1}{3} \chi_i^{-1},$$

where we exploited $\|\mathbf{a}_i^1\|^{-1} \chi_i^{-1} = \|\mathbf{a}_i^1\|^{-1} \min_{2 \leq k \leq d} \|\mathbf{a}_i^k\| \leq \|\mathbf{a}_i^1\| \|\mathbf{a}_i^1\|^{-1} = 1$. Thus, (B3) implies the assumptions (B0), (B1), (B2), (B3), (B4) as well as the hypotheses (H0a), (H0b) and (H1) for all $i = 1, 2, \dots, r$ simultaneously, hereby concluding the main proof.

Part IV: Alternative formulation. It is clear that taking

$$\Delta^T = [\Delta_1^T \ \cdots \ \Delta_r^T] \quad \text{and} \quad \dot{\mathbf{p}}^T = [\text{vec}(\dot{p}_1)^T \ \cdots \ \text{vec}(\dot{p}_r)^T]$$

satisfies $\dot{\mathbf{p}} \sim \mathbf{p}$ and $\|\Delta\| \leq \|\nabla\|$. By assumption, the columns of K form a basis of the kernel of $T_{\mathbf{p}}$. Hence, the Moore–Penrose pseudoinverse of K is $K^\dagger = (K^H K)^{-1} K^H$. As $K K^\dagger$ is a projector onto the column span of K , the claim $\Delta \in \text{span}(K)^\perp$ is equivalent with $K K^\dagger \Delta = 0$. Since

$$K^H \Delta = \sum_{i=1}^r K_i^H \Delta_i = 0,$$

the proof is concluded. \square

As can be understood from Part I in the proof of the Iterated Scaling Lemma, there even exists an iterative algorithm for obtaining the desired factorization. The lemma essentially proves that it always converges for sufficiently small input perturbations ∇ . An implementation of this algorithm is included in the ancillary files accompanying this paper; some numerical experiments illustrating the Iterated Scaling Lemma will be presented in Section 9.1.

A consequence of the proof of Lemma 5 is that Terracini's matrix in the new representatives \dot{p}_i is not very different from Terracini's matrix in the original unperturbed representatives p_i .

Corollary 6. *Let all the assumptions of Lemma 5 hold, and let ∇ , p_i and \dot{p}_i be as in Lemma 5. Let $T_{\mathbf{p}}$ be Terracini's matrix (4) in $\mathbf{p} = \text{vecr}(p_1, \dots, p_r)$, and let $T_{\dot{\mathbf{p}}}$ be Terracini's matrix in $\dot{\mathbf{p}} = \text{vecr}(\dot{p}_1, \dots, \dot{p}_r)$. Then,*

$$T_{\dot{\mathbf{p}}} = T_{\mathbf{p}} \dot{D} = T_{\mathbf{p}}(I + \dot{E}),$$

where the diagonal matrix \dot{D} tends to the identity as $\|\nabla\|$ tends to zero; specifically, there is a constant $C > 0$ such that $\|\dot{E}\|_2 = \|\dot{D} - I\|_2 \leq C\|\nabla\|$.

Proof. Consider the definition of the representative $p_i^{(k)}$ in (9). The limit for $k \rightarrow \infty$ can be written as

$$\begin{aligned} \dot{p}_i &= \lim_{k \rightarrow \infty} p_i^{(k)} = \left(\mathbf{a}_i^1 \cdot \lim_{k \rightarrow \infty} \prod_{j=2}^d (1 - \epsilon_{j,i}^{(k)})^{-1}, \mathbf{a}_i^2 \cdot \lim_{k \rightarrow \infty} (1 - \epsilon_{2,i}^{(k)}), \dots, \mathbf{a}_i^d \cdot \lim_{k \rightarrow \infty} (1 - \epsilon_{d,i}^{(k)}) \right) \\ (23) \quad &= (\gamma_{1,i} \mathbf{a}_i^1, \gamma_{2,i} \mathbf{a}_i^2, \dots, \gamma_{d,i} \mathbf{a}_i^d) \end{aligned}$$

because of (17). Note that the series

$$(24) \quad \gamma_{j,i} = 1 - \epsilon_{j,i}^{(\infty)} = 1 - \sum_{k=1}^{\infty} v_{j,i}^{(k)} > 0, \quad j = 2, 3, \dots, d,$$

are absolutely convergent because the terms in the infinite sequences may be bounded as in (15), and the sum $\chi_i \lambda^{-1} \sum_{k=1}^{\infty} (\lambda \|\nabla\|)^k$ is absolutely convergent because of (B3). Since

$$(25) \quad \gamma_{1,i} = (\gamma_{2,i} \gamma_{3,i} \cdots \gamma_{d,i})^{-1},$$

the above $\gamma_{j,i}$ yield the explicit expressions for the coefficients of \dot{p}_i . Notice that the above furnishes an alternative proof that the sequence of $p_i^{(k)}$'s converges. It follows immediately from the definition of Terracini's matrix in (4) that we can write

$$\dot{T}_i = T_i \cdot \text{diag} \left(I_{n_1} \cdot \prod_{\substack{j=1 \\ j \neq d}}^d \gamma_{j,i}, \dots, I_{n_d} \cdot \prod_{\substack{j=1 \\ j \neq d}}^d \gamma_{j,i} \right) = T_i \cdot \text{diag}(\gamma_{1,i}^{-1} I_{n_1}, \gamma_{2,i}^{-1} I_{n_2}, \dots, \gamma_{d,i}^{-1} I_{n_d}) = T_i \dot{D}_i,$$

so that Terracini's matrix $T_{\dot{\mathbf{p}}}$ in the points \dot{p}_i is given by

$$T_{\dot{\mathbf{p}}} = [\dot{T}_1 \ \cdots \ \dot{T}_r] = T_{\mathbf{p}} \cdot \text{diag}(\dot{D}_1, \dot{D}_2, \dots, \dot{D}_r) = T_{\mathbf{p}} \dot{D}.$$

From the definition of \dot{D} it suffices demonstrating that $|\gamma_{j,i}^{-1} - 1| \leq C\|\nabla\|$. For $j = 2, 3, \dots, d$, we have

$$(26) \quad |\gamma_{j,i}^{-1} - 1| = |(1 - \epsilon_{j,i}^{(\infty)})^{-1} - 1| \leq \sum_{\kappa=1}^{\infty} |\epsilon_{j,i}^{(\infty)}|^\kappa \leq \sum_{\kappa=1}^{\infty} (2\chi_i \|\nabla\|)^\kappa \leq 4\|\nabla\| \max_{1 \leq i \leq r} \chi_i,$$

where in the second equality we used (16) and (B3), and where in the last step we used $2\chi_i\|\nabla_i\| \leq \frac{1}{2}$ which is true because of (B3). The case of $j = 1$ is due to

$$\begin{aligned}
 |\gamma_{1,i}^{-1} - 1| &= \left| -1 + \prod_{j=2}^d (1 - \epsilon_{j,i}^{(\infty)}) \right| = \left| -1 + 1 + \sum_{\kappa=1}^d \sum_{\substack{\|\ell\|_1=\kappa, \\ \ell \leq 1}} (-1)^\kappa \prod_{j=2}^d (\epsilon_{j,i}^{(\infty)})^{\ell_j} \right| \\
 (27) \quad &\leq \sum_{\kappa=1}^{\infty} \sum_{\|\ell\|_1=\kappa} \prod_{j=2}^d |\epsilon_{j,i}^{(\infty)}|^{\ell_j} \leq \sum_{\kappa=1}^{\infty} ((d-1)C'_i)^{\kappa} = \frac{(d-1)C'_i}{1 - (d-1)C'_i} \leq 4(d-1)\|\nabla\| \max_{1 \leq i \leq r} \chi_i;
 \end{aligned}$$

herein, the inequality $\ell \leq 1$ is meant componentwise, the second inequality is by (21), and the last step used (16), (B0) and (B3). Letting $C = 4(d-1) \max_{1 \leq i \leq r} \chi_i$ then concludes the proof. \square

6. A CONDITION NUMBER

The main result of this paper states that the least nonzero singular value of Terracini's matrix is an absolute condition number with respect to the distance measure in Definition 2.

Theorem 7 (Absolute condition number). *Let $N = r(\Sigma+1)$. Let $b_i = (\mathbf{a}_i^1, \mathbf{a}_i^2, \dots, \mathbf{a}_i^d)$ be representatives of $p_i = \text{Seg}(b_i) \in \mathcal{S}_{\mathbb{F}}$, $i = 1, \dots, r$. Let $\mathbf{p} = \text{vecr}(b_1, b_2, \dots, b_r)$ be given. Assume that*

$$\mathfrak{A} = f(\mathbf{p}) = \sum_{i=1}^r \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^d \in \sigma_r^0(\mathcal{S}_{\mathbb{F}}) \subset \mathbb{F}^{\Pi}$$

is robustly r -identifiable, and let ς_N denote the N th largest singular value of Terracini's matrix $T_{\mathbf{p}}$. Then, the absolute condition number of the rank decomposition problem at \mathbf{p} is

$$\kappa_A = \lim_{\epsilon \rightarrow 0} \max_{\substack{\|\Delta \mathfrak{A}\| \leq \epsilon, \\ \mathfrak{A} + \Delta \mathfrak{A} \in \sigma_r^0(\mathcal{S}_{\mathbb{F}})}} \frac{d(\mathbf{p}, f^{\dagger}(\mathfrak{A} + \Delta \mathfrak{A}))}{\|\Delta \mathfrak{A}\|} = \varsigma_N^{-1}.$$

If $\varsigma_N = 0$ or \mathfrak{A} is not robustly r -identifiable, then the condition number is defined to be ∞ .

Proof. Observe that f is an analytic multivariate polynomial that is homogeneous of degree d in the $M = r(\Sigma + d)$ variables x_i , so that it has a finite, convergent Taylor series expansion at every point. Let $\Delta \mathbf{p} \in \mathbb{F}^M$ be arbitrary. For $f_i(\mathbf{p} + \Delta \mathbf{p})$, where f_i is the i th component of the vector function f , this expansion around \mathbf{p} is given by

$$f_i(\mathbf{p} + \Delta \mathbf{p}) = f_i(\mathbf{p}) + \sum_{1 \leq k_1 \leq M} \Delta p_{k_1} \left(\frac{\partial f_i}{\partial x_{k_1}} \right)(\mathbf{p}) + \sum_{1 \leq k_1 < k_2 \leq M} \Delta p_{k_1} \Delta p_{k_2} \left(\frac{\partial^2 f_i}{\partial x_{k_1} \partial x_{k_2}} \right)(\mathbf{p}) + \dots,$$

where we exploited the observation that the terms corresponding to the higher-order partial derivatives with respect to the same variable are zero, so that we can get rid of the multinomial coefficients. Then,

$$f(\mathbf{p} + \Delta \mathbf{p}) = f(\mathbf{p}) + T_{\mathbf{p}} \Delta \mathbf{p} + \mathcal{O}(\|\Delta \mathbf{p}\|^2),$$

where $T_{\mathbf{p}}$ is Terracini's matrix as in (4).

Let \mathcal{N} be the neighborhood of \mathfrak{A} where all tensors are robustly r -identifiable and hence of rank r . Let $\mathfrak{A}' \in \mathcal{N}$ be arbitrary. Then, there exist several $\Delta \mathbf{p}'$ such that $\mathfrak{A}' = \mathfrak{A} + \Delta \mathfrak{A} = f(\mathbf{p} + \Delta \mathbf{p}')$. The remainder of the proof will be slightly easier if we enforce one unique particular choice. Note that by r -identifiability, $f^{\dagger}(\mathfrak{A}') = \mathcal{T}(\mathbf{p} + \Delta \mathbf{p}')$ so that the set of rank-1 tensors appearing in the decomposition of \mathfrak{A}' is unique. The r elements of this set can be ordered uniquely with respect to the lexicographic total order \leq on \mathbb{F}^{Π} .⁷ This removes the permutation ambiguity, leaving only the scaling indeterminacies. These can be removed as follows. Let $b'_i = (\mathbf{x}_i^1, \mathbf{x}_i^2, \dots, \mathbf{x}_i^d)$ be a set of representatives such that the corresponding rank-1 tensors $\text{Seg}(b'_i)$ are sorted. Then, the following set is uniquely defined regardless of the particular choice of representatives for the rank-1 tensors:

$$(28) \quad \mathcal{C} = \left\{ j_{k,i} = \min \left\{ \arg \max_{1 \leq \ell \leq n_k} |x_{\ell,i}^{(k)}| \right\} \mid k = 2, 3, \dots, d, i = 1, 2, \dots, r \right\},$$

⁷For $\mathbb{F} = \mathbb{C} \simeq \mathbb{R}^2$, one should consider the lexicographic order after the identification $\mathbb{C}^{\Pi} = \mathbb{R}^{2\Pi}$.

where $x_{\ell,i}^{(k)}$ is the ℓ th element of $\mathbf{x}_i^k \in \mathbb{F}^{n_k}$. Note that $|x_{j_k,i,i}^{(k)}| = \|\mathbf{x}_i^k\|_\infty > 0$ by r -identifiability, and that $|\mathcal{C}| = r(d-1)$. It follows from r -identifiability that there is just one choice of vectorized factor matrices $\mathbf{p} + \Delta\mathbf{p} = \text{vecr}(b_1'', b_2'', \dots, b_r'')$ with

$$(29) \quad b_i'' = (\mathbf{b}_i^1, \mathbf{b}_i^2, \dots, \mathbf{b}_i^d), \quad \text{Seg}(b_1'') < \text{Seg}(b_2'') < \dots < \text{Seg}(b_r''), \quad \text{and} \quad \forall j_{k,i} \in \mathcal{C} : b_{j_{k,i},i}^{(k)} = 1$$

such that $\mathfrak{A}' = f(\mathbf{p} + \Delta\mathbf{p})$, as the last set of equations completely fixes the scaling indeterminacies. Note that the order of the rank-1 tensors is strict, for otherwise the rank of \mathfrak{A}' would be strictly less than r . Since \mathbf{p} is fixed, $\Delta\mathbf{p}$ is uniquely determined by the foregoing procedure.

Write $\Delta\mathfrak{A} = f(\mathbf{p} + \Delta\mathbf{p}) - f(\mathbf{p})$. Since r -identifiability holds in \mathcal{N} —but only up to the action of \mathcal{T} —it follows that $f^\dagger(\mathfrak{A}') = \mathcal{T}(\mathbf{p} + \Delta\mathbf{p})$. Choose any

$$(30) \quad \widehat{S} \in \mathcal{Z} = \arg \min_{\widetilde{S} \in \mathcal{T}} \|\mathbf{p} - S(\mathbf{p} + \Delta\mathbf{p})\|;$$

note that \widehat{S} depends on \mathfrak{A}' , however I will not make such dependencies explicit in the notation for the sake of brevity. We can write $\widehat{S}(\mathbf{p} + \Delta\mathbf{p}) = \mathbf{p} + \widetilde{\Delta\mathbf{p}}$, so that for every $\mathfrak{A}' \in \mathcal{N}$ the following holds:

$$\mathfrak{A}' = f(\mathbf{p} + \Delta\mathbf{p}) = f(\mathbf{p} + \widetilde{\Delta\mathbf{p}}).$$

Notice that

$$\|\widetilde{\Delta\mathbf{p}}\| = d(\mathbf{p}, f^\dagger(\mathfrak{A} + \Delta\mathfrak{A})),$$

so we can just analyze this norm. It is important to note that while $\widetilde{\Delta\mathbf{p}}$ depends on the choice of \widehat{S} , the norm of this vector *is independent* of this choice.

In the remainder of the proof, most considered quantities, such as $\widetilde{\Delta\mathbf{p}}$, depend on $\Delta\mathfrak{A}$, which is arbitrary, and the choice of \widehat{S} (which, in turn, depends on $\Delta\mathfrak{A}$). To be very precise, one could indicate these dependencies in the notation by writing, for example, $\widetilde{\Delta\mathbf{p}}(\Delta\mathfrak{A}, \widehat{S}(\Delta\mathfrak{A}))$, but I will refrain from doing so wherever no confusion may arise for avoiding the obvious notational burden.

Part I: Continuity. First, we prove that

$$(31) \quad \lim_{\substack{\|\Delta\mathfrak{A}\| \rightarrow 0, \\ \mathfrak{A} + \Delta\mathfrak{A} \in \sigma_r^0(\mathcal{S}_{\mathbb{F}})}} \|\widetilde{\Delta\mathbf{p}}(\Delta\mathfrak{A})\| \rightarrow 0$$

by eliminating $r(d-1)$ parameters as follows. Let $\mathbf{x} = \text{vecr}(\chi_1, \chi_2, \dots, \chi_r) \in \mathbb{F}^{r(\Sigma+d)}$ be variables, where $\chi_i = (\mathbf{x}_i^1, \mathbf{x}_i^2, \dots, \mathbf{x}_i^d)$. We will represent the tensor $\mathfrak{X} = f(\mathbf{x})$ using the expected number of parameters N —provided that $\mathfrak{X} \in \mathcal{N}$ is close to \mathfrak{A} —by choosing vectorized factor matrices that lie in the subspace⁸ of $\mathbb{F}^{r(\Sigma+d)}$ given by the system of $r(d-1)$ equations

$$U_{\mathcal{C}, \mathbf{p}} = \{x_{j_{k,i},i}^{(k)} = a_{j_{k,i},i}^{(k)} \mid j_{k,i} \in \mathcal{C}\},$$

where \mathcal{C} is as in (28), and $x_{\ell,i}^{(k)}$ and $a_{\ell,i}^{(k)}$ denote the ℓ th element of \mathbf{x}_i^k and \mathbf{a}_i^k respectively. The change of variables

$$\eta : \mathbf{x} \mapsto \dot{\mathbf{x}} = \text{vecr}(\dot{\chi}_1, \dot{\chi}_2, \dots, \dot{\chi}_r),$$

where $\dot{\chi}_i = (\dot{\mathbf{x}}_i^1, \dot{\mathbf{x}}_i^2, \dots, \dot{\mathbf{x}}_i^d)$, $\dot{\mathbf{x}}_i^k = [\dot{x}_{1,i}^{(k)} \dots \dot{x}_{n_k,i}^{(k)}]^T$, and

$$(32a) \quad \dot{x}_{\ell,i}^{(1)} = x_{\ell,i}^{(1)} \prod_{k=2}^d \frac{x_{j_{k,i},i}^{(k)}}{a_{j_{k,i},i}^{(k)}}, \quad \ell = 1, 2, \dots, n_1, \quad i = 1, 2, \dots, r,$$

$$(32b) \quad \dot{x}_{\ell,i}^{(k)} = a_{j_{k,i},i}^{(k)} \frac{x_{\ell,i}^{(k)}}{x_{j_{k,i},i}^{(k)}}, \quad \ell = 1, 2, \dots, n_k, \quad i = 1, 2, \dots, r, \quad k = 2, 3, \dots, d,$$

allows one to represent every tensor $f(\mathbf{x})$ with $x_{j_{k,i},i}^{(k)} \neq 0$, $j_{k,i} \in \mathcal{C}$ on $U_{\mathcal{C}, \mathbf{p}}$ as $\dot{\mathbf{x}}$. In particular, we may assume that all tensors $\mathfrak{A}' \in \mathcal{N}$ can be represented on $U_{\mathcal{C}, \mathbf{p}}$. Otherwise, one should shrink the open neighborhood \mathcal{N} until $\mathcal{N} \subset f(U_{\mathcal{C}, \mathbf{p}})$ is satisfied; this is always possible due to the particular choice of \mathcal{C} which ensures that $|a_{j_{k,i},i}^{(k)}| = \|\mathbf{a}_i^k\|_\infty > 0$ if $j_{k,i} \in \mathcal{C}$ —the ∞ -norm is nonzero because \mathfrak{A} is robustly

⁸This can be thought of as a normalized representation of the vectorized factor matrices.

r -identifiable and hence of rank r . Note that the definition of η is such that it simply chooses an alternative representative for each of the rank-1 tensors χ_i , so that

$$\exists \bar{S} \in \mathcal{T} : \eta(\mathbf{x}) = \bar{S}\mathbf{x};$$

note that \bar{S} depends on \mathbf{x} , but I will not indicate this in the notation. Define for $k = 2, 3, \dots, d$ and $i = 1, 2, \dots, r$, the functions

$$\begin{aligned} \pi_{j_{k,i}, \mathfrak{A}} : \mathbb{F}^{n_k} &\rightarrow \mathbb{F}^{n_k-1}, \mathbf{z} \mapsto [z_1 \quad \cdots \quad z_{j_{k,i}-1} \quad z_{j_{k,i}+1} \quad \cdots \quad z_{n_k}]^T, \\ \pi_{j_{k,i}, \mathfrak{A}}^{-1} : \mathbb{F}^{n_k-1} &\rightarrow \mathbb{F}^{n_k}, \mathbf{z} \mapsto [z_1 \quad \cdots \quad z_{j_{k,i}-1} \quad a_{j_{k,i}}^{(k)} \quad z_{j_{k,i}} \quad \cdots \quad z_{n_k-1}]^T, \end{aligned}$$

where $j_{k,i} \in \mathcal{C}$. Define also

$$\begin{aligned} \pi_{\mathfrak{A}} : \mathbb{F}^M &\rightarrow \mathbb{F}^N & \pi_{\mathfrak{A}}^{-1} : \mathbb{F}^N &\rightarrow \mathbb{F}^M \\ \begin{bmatrix} \mathbf{z}_i^1 \\ \mathbf{z}_i^2 \\ \vdots \\ \mathbf{z}_i^d \end{bmatrix}_{i=1,\dots,r} &\mapsto \begin{bmatrix} \mathbf{z}_i^1 \\ \pi_{j_{2,i}, \mathfrak{A}}(\mathbf{z}_i^2) \\ \vdots \\ \pi_{j_{d,i}, \mathfrak{A}}(\mathbf{z}_i^d) \end{bmatrix}_{i=1,\dots,r}, & \text{and} & \begin{bmatrix} \mathbf{z}_i^1 \\ \mathbf{z}_i^2 \\ \vdots \\ \mathbf{z}_i^d \end{bmatrix}_{i=1,\dots,r} &\mapsto \begin{bmatrix} \mathbf{z}_i^1 \\ \pi_{j_{2,i}, \mathfrak{A}}^{-1}(\mathbf{z}_i^2) \\ \vdots \\ \pi_{j_{d,i}, \mathfrak{A}}^{-1}(\mathbf{z}_i^d) \end{bmatrix}_{i=1,\dots,r}. \end{aligned}$$

With the definition of these analytic functions in place, we define

$$\begin{aligned} f_{\mathfrak{A}} : \mathbb{F}^N &\rightarrow \mathbb{F}^{\Pi} \\ \mathbf{z} &\mapsto (f \circ \pi_{\mathfrak{A}}^{-1})(\mathbf{z}). \end{aligned}$$

Now it follows that for an arbitrary choice of $\mathfrak{A}' = f(\mathbf{p} + \Delta\mathbf{p}) \in \mathcal{N}$ there exists an $\bar{S} \in \mathcal{T}$ such that

$$\mathfrak{A}' = \mathfrak{A} + \Delta\mathfrak{A} = f(\mathbf{p} + \Delta\mathbf{p}) = f_{\mathfrak{A}}(\pi_{\mathfrak{A}}(\eta(\mathbf{p} + \Delta\mathbf{p}))) = f_{\mathfrak{A}}(\pi_{\mathfrak{A}}(\bar{S}(\mathbf{p} + \Delta\mathbf{p}))).$$

Let $\bar{\mathbf{p}} = \pi_{\mathfrak{A}}(\eta(\mathbf{p})) = \pi_{\mathfrak{A}}(\mathbf{p})$ and $\bar{\Delta\mathbf{p}} = \pi_{\mathfrak{A}}(\bar{S}(\mathbf{p} + \Delta\mathbf{p})) - \bar{\mathbf{p}}$. As $f_{\mathfrak{A}}$ is an analytic function over \mathbb{F} in N variables, it is continuously differentiable with Jacobian \bar{J} at $\bar{\mathbf{p}}$. Its rank is the maximum value N . This can be understood by observing that the Jacobian of $f \circ \pi_{\mathfrak{A}}^{-1}$ is—by the chain rule—given by

$$\bar{J} = J_{f \circ \pi_{\mathfrak{A}}^{-1}}(\bar{\mathbf{p}}) = J_f(\pi_{\mathfrak{A}}^{-1}(\bar{\mathbf{p}}))J_{\pi_{\mathfrak{A}}^{-1}}(\bar{\mathbf{p}}) = J_f(\mathbf{p})\hat{I} = T_{\mathbf{p}}\hat{I},$$

where $\hat{I} \in \mathbb{R}^{r(\Sigma+d) \times r(\Sigma+1)}$ is an identity matrix from which certain columns have been removed. The implication is that the Jacobian \bar{J} is formed by a subset of the columns of f 's Jacobian, i.e., Terracini's matrix $T_{\mathbf{p}}$. In fact, from the definition of $\pi_{\mathfrak{A}}^{-1}$ one verifies that \bar{J} lacks the columns

$$(33) \quad \mathbf{a}_i^1 \otimes \cdots \otimes \mathbf{a}_i^{k-1} \otimes \mathbf{e}_{j_{k,i}} \otimes \mathbf{a}_i^{k+1} \otimes \cdots \otimes \mathbf{a}_i^d, \quad j_{k,i} \in \mathcal{C},$$

relative to $T_{\mathbf{p}}$. I claim that these columns in (33) are linearly dependent on the other columns of $T_{\mathbf{p}}$. Indeed, fix $k = 2, 3, \dots, d$ and $i = 1, 2, \dots, r$, and then

$$T_{\mathbf{p}}(\mathbf{k}_i^k \otimes \mathbf{e}_i) = \sum_{j=1}^{n_1} a_{j,i}^{(1)} \mathbf{e}_j \otimes \mathbf{a}_i^2 \otimes \cdots \otimes \mathbf{a}_i^d - \sum_{j=1}^{n_k} a_{j,i}^{(k)} \mathbf{a}_i^1 \otimes \cdots \otimes \mathbf{a}_i^{k-1} \otimes \mathbf{e}_j \otimes \mathbf{a}_i^{k+1} \otimes \cdots \otimes \mathbf{a}_i^d = 0.$$

As $a_{j_{k,i}}^{(k)} \neq 0$, $j_{k,i} \in \mathcal{C}$, by construction of \mathcal{C} , it follows that (33) is linearly dependent on the other columns in the above expression, and may hence be removed from $T_{\mathbf{p}}$ without affecting its range. By inspection it is clear that this procedure drops exactly $r(d-1)$ linearly independent columns of $T_{\mathbf{p}}$, resulting in \bar{J} whose column span coincides with the column span of the original Terracini's matrix. As the latter had dimension N ,⁹ it follows that $\bar{J} \in \mathbb{F}^{\Pi \times N}$ has linearly independent columns.

I claim that $f_{\mathfrak{A}}$ is a local diffeomorphism. By our assumption on robust r -identifiability, \mathcal{N} is an open neighborhood of a smooth manifold of dimension N . Hence, there exists a chart (U, ϕ^{-1}) where $U \subset \mathcal{N}$ is an open neighborhood of \mathfrak{A} and a local diffeomorphism ϕ^{-1} between U and an open neighborhood $V = \phi^{-1}(U) \subset \mathbb{F}^N$. Restrict \mathcal{N} to the open neighborhood $U \subset \mathcal{N}$. Let $\mathfrak{A} = \phi(\mathbf{q}) \in \mathcal{N}$. Since ϕ is a local diffeomorphism, $\phi^{-1} \circ \phi|_V = \text{Id}_{\mathbb{F}^N}$. Hence, the Jacobian matrix at $\mathbf{q} \in V$ of the

⁹The condition number is defined to be ∞ if this rank condition is not satisfied.

composition $\phi^{-1} \circ \phi|_V$, namely $J_{\phi^{-1}}(\phi(\mathbf{q}))J_{\phi}(\mathbf{q})$ by the chain rule, is the $N \times N$ identity matrix. Since $J_{\phi}(\mathbf{q}) \in \mathbb{F}^{\Pi \times N}$ and $J_{\phi^{-1}}(\phi(\mathbf{q})) \in \mathbb{F}^{N \times \Pi}$ with $\Pi \geq N$, it follows that both matrices are of maximal rank N . By definition, $J_{\phi}(\mathbf{q})$ is contained in the tangent space $T_{\mathfrak{A}\sigma_r(\mathcal{S}_{\mathbb{F}})}$. By our assumption on robust r -identifiability, $\text{span}(\overline{J}) = \text{span}(T_{\mathbf{p}}) = T_{\mathfrak{A}\sigma_r(\mathcal{S}_{\mathbb{F}})}$. Since $\dim \text{span}(J_{\phi}(\mathbf{q})) = N = \dim \text{span}(\overline{J})$ it follows that $J_{\phi}(\mathbf{q})Z = \overline{J}$ for some nonsingular matrix $Z \in \mathbb{F}^{N \times N}$. Hence, the Jacobian matrix of the composition $g : \phi^{-1} \circ f_{\mathfrak{A}}$ at $\overline{\mathbf{p}}$ is given by $J_{\phi^{-1}}(f_{\mathfrak{A}}(\overline{\mathbf{p}}))\overline{J} = J_{\phi^{-1}}(\phi(\mathbf{q}))J_{\phi}(\mathbf{q})Z = Z$, where the first equality is because $\mathfrak{A} = f_{\mathfrak{A}}(\overline{\mathbf{p}}) = \phi(\mathbf{q})$. Now, the smooth function $g : \mathbb{F}^N \rightarrow \mathbb{F}^N$ has nonsingular Jacobian matrix Z , hence by the Inverse Function Theorem [39, Theorem 1A.1] it has a smooth inverse function g^{-1} . Thus, g is a local diffeomorphism between a neighborhood of $\mathbf{p} \in \mathbb{F}^N$ and a neighborhood of $\mathbf{q} \in \mathbb{F}^N$. Letting now $f_{\mathfrak{A}}^{-1}$ be the composition $g^{-1} \circ \phi^{-1}$, which is a composition of local diffeomorphisms, it follows that $f_{\mathfrak{A}}$ is a local diffeomorphism between an open neighborhood of $\mathbf{p} \in \mathbb{F}^N$ and an open neighborhood of $\mathfrak{A} \in \mathcal{N}$. Further restrict \mathcal{N} to the neighborhood where this local diffeomorphism is defined.

As $f_{\mathfrak{A}}^{-1}$ is locally smooth, it follows immediately that $\Delta \mathfrak{A} \rightarrow 0$ implies $\overline{\Delta \mathbf{p}} \rightarrow 0$. Both \mathbf{p} and $\overline{S}(\mathbf{p} + \Delta \mathbf{p})$ live in $U_{\mathcal{C}, \mathbf{p}}$, so that their values at all positions $j_{k,i} \in \mathcal{C}$ agree. As $\pi_{\mathfrak{A}}^{-1}(\mathbf{x})$ is analytic, it follows that $\overline{\Delta \mathbf{p}} \rightarrow 0$ implies $\overline{S}(\mathbf{p} + \Delta \mathbf{p}) - \mathbf{p} \rightarrow 0$ as well. Since

$$\|\widetilde{\Delta \mathbf{p}}\| = \min_{S \in \mathcal{T}} \|\mathbf{p} - S(\mathbf{p} + \Delta \mathbf{p})\| \leq \|\mathbf{p} - \overline{S}(\mathbf{p} + \Delta \mathbf{p})\|,$$

equation (31) follows.

Part II: Sandwiching. For convenience, let us denote the set

$$\mathcal{G}_{\epsilon} = \{\Delta \mathfrak{A} \in \mathbb{F}^{\Pi} \setminus \{0\} \mid \|\Delta \mathfrak{A}\| \leq \epsilon \text{ and } \mathfrak{A} + \Delta \mathfrak{A} \in \sigma_r^0(\mathcal{S}_{\mathbb{F}})\}.$$

Then, the second part consists of demonstrating that

$$\lim_{\epsilon \rightarrow 0} \max_{\Delta \mathfrak{A} \in \mathcal{G}_{\epsilon}} \frac{\|\widetilde{\Delta \mathbf{p}}(\Delta \mathfrak{A})\|}{\|\Delta \mathfrak{A}\|} = \lim_{\epsilon \rightarrow 0} \max_{\Delta \mathfrak{A} \in \mathcal{G}_{\epsilon}} \frac{\|T_{\mathbf{p}} \mathbf{x}(\Delta \mathfrak{A})\|}{\|\mathbf{x}(\Delta \mathfrak{A})\|} = (\varsigma_N(T_{\mathbf{p}}))^{-1}$$

whereby $\mathbf{x}(\Delta \mathfrak{A}) \in \mathbb{F}^{r(\Sigma+d)}$. The key consists of exploiting the Iterated Scaling Lemma for transforming any $\widetilde{\Delta \mathbf{p}}$ into a new vector $\mathbf{x}(\Delta \mathfrak{A})$ that is orthogonal to $T_{\mathbf{p}}$'s kernel.

Since Terracini's matrix $T_{\mathbf{p}}$ has a kernel $\mathcal{K} = \text{span}(K)$ of dimension exactly equal to $r(d-1)$, we may factorize $\widetilde{\Delta \mathbf{p}}$ as

$$\widetilde{\Delta \mathbf{p}} = \Delta + \nabla, \quad \text{where } \Delta \in \mathcal{K}^{\perp} \text{ and } \nabla \in \mathcal{K}.$$

Notice that $\langle \nabla, \Delta \rangle = 0$, so that $\|\nabla\| \leq \|\widetilde{\Delta \mathbf{p}}\|$ and $\|\Delta\| \leq \|\widetilde{\Delta \mathbf{p}}\|$. From Lemma 5 it follows that for sufficiently small $\|\Delta \mathfrak{A}\| > 0$, and, hence, for sufficiently small $\|\widetilde{\Delta \mathbf{p}}\| > 0$, we have that

$$\mathbf{p} + \nabla = \dot{\mathbf{p}} + \Delta' \quad \text{with } \dot{\mathbf{p}} \sim \mathbf{p}, \Delta' \in \mathcal{K}^{\perp}, \text{ and } \|\Delta'\| \leq \|\nabla\|.$$

Let $\dot{\Delta} = \Delta' + \Delta \in \mathcal{K}^{\perp}$. As particular consequences, we have that $f(\mathbf{p}) = f(\dot{\mathbf{p}})$ and $f(\mathbf{p} + \widetilde{\Delta \mathbf{p}}) = f(\dot{\mathbf{p}} + \dot{\Delta})$. The Taylor series expansion about the vectorized factor matrices $\dot{\mathbf{p}}$ is

$$f(\mathbf{p} + \widetilde{\Delta \mathbf{p}}) = f(\dot{\mathbf{p}} + \dot{\Delta}) = f(\dot{\mathbf{p}}) + T_{\dot{\mathbf{p}}} \dot{\Delta} + \mathcal{O}(\|\dot{\Delta}\|^2) = f(\mathbf{p}) + T_{\mathbf{p}} \dot{\Delta} + \mathcal{O}(\|\dot{\Delta}\|^2);$$

thus, there exist a constant $C > 0$ such that for all sufficiently small $\widetilde{\Delta \mathbf{p}}$ it holds that

$$\|T_{\dot{\mathbf{p}}} \dot{\Delta}\| - C\|\dot{\Delta}\|^2 \leq \|f(\mathbf{p} + \widetilde{\Delta \mathbf{p}}) - f(\mathbf{p})\| \leq \|T_{\dot{\mathbf{p}}} \dot{\Delta}\| + C\|\dot{\Delta}\|^2.$$

It follows from the Iterated Scaling Lemma that if $f(\mathbf{p} + \widetilde{\Delta \mathbf{p}}) \neq f(\mathbf{p})$ then $\|\dot{\Delta}\| \neq 0$.¹⁰ Indeed, the contrapositive $\|\dot{\Delta}\| = 0$ yields $\dot{\Delta} = \mathbf{p} + \widetilde{\Delta \mathbf{p}} - \dot{\mathbf{p}} = 0$, so that $f(\mathbf{p} + \widetilde{\Delta \mathbf{p}}) = f(\dot{\mathbf{p}}) = f(\mathbf{p})$. Thus, if $\|\Delta \mathfrak{A}\| > 0$ then $\|\dot{\Delta}\| > 0$, so that we may write

$$-C\|\dot{\Delta}\| \leq \frac{\|f(\mathbf{p} + \widetilde{\Delta \mathbf{p}}) - f(\mathbf{p})\|}{\|\dot{\Delta}\|} - \frac{\|T_{\dot{\mathbf{p}}} \dot{\Delta}\|}{\|\dot{\Delta}\|} \leq C\|\dot{\Delta}\|.$$

¹⁰At this point where we invoke separability, the argument for the premetric $\widehat{d}(\cdot, \cdot)$ fails.

Hence, for sufficiently small ϵ and $\|\dot{\Delta}\|$, we obtain the useful bound

$$(34) \quad \left| \frac{\|\Delta\mathfrak{A}\|}{\|\dot{\Delta}\|} - \frac{\|T_{\mathbf{p}}\dot{\Delta}\|}{\|\dot{\Delta}\|} \right| \leq C \cdot \|\dot{\Delta}\| \leq C \cdot \max_{\Delta\mathfrak{A} \in \mathcal{G}_\epsilon} \max_{\widehat{S}(\Delta\mathfrak{A}) \in \mathcal{Z}} \|\dot{\Delta}(\Delta\mathfrak{A}, \widehat{S}(\Delta\mathfrak{A}))\|.$$

Notice that all of the dotted vectors ultimately depend on $\Delta\mathfrak{A}$. Consider now

$$\frac{\|T_{\mathbf{p}}\dot{\Delta}\|}{\|\dot{\Delta}\|} - \frac{\|T_{\mathbf{p}}\dot{E}\dot{\Delta}\|}{\|\dot{\Delta}\|} \leq \frac{\|T_{\mathbf{p}}\dot{\Delta}\|}{\|\dot{\Delta}\|} = \frac{\|T_{\mathbf{p}}\dot{\Delta} + T_{\mathbf{p}}\dot{E}\dot{\Delta}\|}{\|\dot{\Delta}\|} \leq \frac{\|T_{\mathbf{p}}\dot{\Delta}\|}{\|\dot{\Delta}\|} + \frac{\|T_{\mathbf{p}}\dot{E}\dot{\Delta}\|}{\|\dot{\Delta}\|},$$

where \dot{E} is as in Corollary 6. Then, we similarly have for sufficiently small ϵ that

$$(35) \quad \left| \frac{\|T_{\mathbf{p}}\dot{\Delta}\|}{\|\dot{\Delta}\|} - \frac{\|T_{\mathbf{p}}\dot{E}\dot{\Delta}\|}{\|\dot{\Delta}\|} \right| \leq \frac{\|T_{\mathbf{p}}\dot{E}\dot{\Delta}\|}{\|\dot{\Delta}\|} \leq \|T_{\mathbf{p}}\|_2 \|\dot{E}\|_2 \leq \|T_{\mathbf{p}}\|_2 \max_{\Delta\mathfrak{A} \in \mathcal{G}_\epsilon} \max_{\widehat{S}(\Delta\mathfrak{A}) \in \mathcal{Z}} \|\dot{E}(\Delta\mathfrak{A}, \widehat{S}(\Delta\mathfrak{A}))\|_2.$$

Combining (34) and (35) allows us to infer that

$$(36a) \quad \frac{\|\Delta\mathfrak{A}\|}{\|\dot{\Delta}\|} = \frac{\|T_{\mathbf{p}}\dot{\Delta}\|}{\|\dot{\Delta}\|} + \varepsilon,$$

where

$$(36b) \quad |\varepsilon| \leq \vartheta = \left(C \cdot \max_{\Delta\mathfrak{A} \in \mathcal{G}_\epsilon} \max_{\widehat{S}(\Delta\mathfrak{A}) \in \mathcal{Z}} \|\dot{\Delta}(\Delta\mathfrak{A}, \widehat{S}(\Delta\mathfrak{A}))\| \right) + (\|T_{\mathbf{p}}\|_2 \cdot \max_{\Delta\mathfrak{A} \in \mathcal{G}_\epsilon} \max_{\widehat{S}(\Delta\mathfrak{A}) \in \mathcal{Z}} \|\dot{E}(\Delta\mathfrak{A}, \widehat{S}(\Delta\mathfrak{A}))\|_2).$$

It remains to relate this expression to $\|\widetilde{\Delta\mathbf{p}}\|/\|\Delta\mathfrak{A}\|$.

An upper bound on $\|\widetilde{\Delta\mathbf{p}}\|$ in terms of $\dot{\Delta}$ is obtained by observing that

$$\widehat{S}(\mathbf{p} + \Delta\mathbf{p}) = \mathbf{p} + \widetilde{\Delta\mathbf{p}} = \dot{\mathbf{p}} + \dot{\Delta} = S\mathbf{p} + \dot{\Delta}$$

for some invertible $S \in \mathcal{T}$. It follows that

$$\|\widetilde{\Delta\mathbf{p}}\| = \min_{D \in \mathcal{T}} \|D(\mathbf{p} + \Delta\mathbf{p}) - \mathbf{p}\| \leq \|S^{-1}\widehat{S}(\mathbf{p} + \Delta\mathbf{p}) - \mathbf{p}\| = \|S^{-1}\dot{\Delta}\| \leq \|S^{-1}\|_2 \|\dot{\Delta}\|,$$

where the first inequality arises because \mathcal{T} is a multiplicative group. As $S^{-1} \in \mathcal{T}$ so $S \in \mathcal{T}$, we can write $S = PD$ where $P \in \mathcal{P}$ and $D \in \mathcal{B}$, so that $\|S^{-1}\|_2 = \|D^{-1}\|_2$. It follows from the Iterated Scaling Lemma that the diagonal entries of D are the $\gamma_{j,i}$'s in (23). Then, it follows immediately from (26) and (27) that the largest diagonal entry in absolute value of D^{-1} is smaller than $1 + 4(d-1)\|\widetilde{\Delta\mathbf{p}}\| \cdot \max_i \chi_i = 1 + \delta$, where χ_i is as in (14). It follows from Part I that $\delta \rightarrow 0$ as $\epsilon \rightarrow 0$. A lower bound on $\|\widetilde{\Delta\mathbf{p}}\|$ in terms of $\dot{\Delta}$ is obtained as follows. From the triangle inequality, $\|\dot{\Delta}\| \leq \|\Delta'\| + \|\Delta\|$ so that

$$\begin{aligned} \|\dot{\Delta}\|^2 &\leq \|\Delta'\|^2 + 2\|\Delta'\|\|\Delta\| + \|\Delta\|^2 \\ &= \|\Delta'\|(\|\Delta'\| + 2\|\Delta\|) + \|\Delta\|^2 \leq \|\nabla\|^2(4\lambda^2\|\nabla\|^2 + 4\lambda\|\Delta\|) + \|\Delta\|^2, \end{aligned}$$

where the last step is due to Lemma 5. Note that λ is a constant. Provided that $\|\Delta\mathfrak{A}\|$ and, hence, $\|\widetilde{\Delta\mathbf{p}}\|$, are sufficiently small, i.e., so that the inequality

$$4\lambda^2\|\nabla\|^2 + 4\lambda\|\Delta\| \leq 4\lambda^2\|\widetilde{\Delta\mathbf{p}}\|^2 + 4\lambda\|\widetilde{\Delta\mathbf{p}}\| \leq 1$$

is satisfied, then the foregoing implies the lower bound

$$\|\dot{\Delta}\|^2 \leq \|\nabla\|^2 + \|\Delta\|^2 = \|\nabla\|^2 + 2 \cdot \text{Re}\langle \nabla, \Delta \rangle + \|\Delta\|^2 = \|\widetilde{\Delta\mathbf{p}}\|^2,$$

where the first equality is because of the orthogonality in the Hermitian inner product of ∇ and Δ . Thus, by combining the lower and upper bound and dividing by $\|\Delta\mathfrak{A}\|$, we can conclude that for every nonzero $\Delta\mathfrak{A}$ of sufficiently small norm the following relations hold true

$$(37) \quad \frac{\|\dot{\Delta}(\Delta\mathfrak{A}, \widehat{S}(\Delta\mathfrak{A}))\|}{\|\Delta\mathfrak{A}\|} \leq \frac{\|\widetilde{\Delta\mathbf{p}}(\Delta\mathfrak{A})\|}{\|\Delta\mathfrak{A}\|} \leq (1 + \delta(\Delta\mathfrak{A})) \frac{\|\dot{\Delta}(\Delta\mathfrak{A}, \widehat{S}(\Delta\mathfrak{A}))\|}{\|\Delta\mathfrak{A}\|},$$

provided that $\epsilon \geq \|\Delta \mathfrak{A}\|$ is small enough and where I explicitly indicated the dependence on $\Delta \mathfrak{A}$ and \widehat{S} , which in turn depends on the former. Let

$$\Delta \mathfrak{A}' \in \arg \max_{\Delta \mathfrak{A} \in \mathcal{G}_\epsilon} \max_{\widehat{S}(\Delta \mathfrak{A}) \in \mathcal{Z}} \frac{\|\dot{\Delta}(\Delta \mathfrak{A}, \widehat{S}(\Delta \mathfrak{A}))\|}{\|\Delta \mathfrak{A}\|} \quad \text{and} \quad \Delta \mathfrak{A}'' \in \arg \max_{\Delta \mathfrak{A} \in \mathcal{G}_\epsilon} \frac{\|\widetilde{\Delta \mathbf{p}}(\Delta \mathfrak{A})\|}{\|\Delta \mathfrak{A}\|},$$

where \mathcal{Z} is as in (30). Also let

$$\widehat{S}' \in \arg \max_{\widehat{S} \in \mathcal{Z}} \frac{\|\dot{\Delta}(\Delta \mathfrak{A}', \widehat{S})\|}{\|\Delta \mathfrak{A}'\|}.$$

It follows that

$$(38) \quad \frac{\|\dot{\Delta}(\Delta \mathfrak{A}', \widehat{S}')\|}{\|\Delta \mathfrak{A}'\|} \leq \frac{\|\widetilde{\Delta \mathbf{p}}(\Delta \mathfrak{A}')\|}{\|\Delta \mathfrak{A}'\|} \leq \frac{\|\widetilde{\Delta \mathbf{p}}(\Delta \mathfrak{A}'')\|}{\|\Delta \mathfrak{A}''\|} \leq (1 + \delta(\Delta \mathfrak{A}'')) \frac{\|\dot{\Delta}(\Delta \mathfrak{A}'', \bullet)\|}{\|\Delta \mathfrak{A}''\|} \\ \leq \frac{\|\dot{\Delta}(\Delta \mathfrak{A}', \widehat{S}')\|}{\|\Delta \mathfrak{A}'\|} \cdot \max_{\Delta \mathfrak{A} \in \mathcal{G}_\epsilon} (1 + \delta(\Delta \mathfrak{A})),$$

where the first step is by the first inequality in (37), the second inequality is by optimality of $\Delta \mathfrak{A}''$, the third inequality is because of the second inequality in (37) with the “ \bullet ” indicating that the inequality is valid for all $S(\Delta \mathfrak{A}'') \in \mathcal{Z}$ (so in particular the maximum), and the last step is because of the optimality of $\Delta \mathfrak{A}'$. For continuing, I will drop the explicit notation indicating the dependence on $\Delta \mathfrak{A}$ and \widehat{S} again and we will consider the sequences

$$c_n = \min_{\Delta \mathfrak{A}_n \in \mathcal{G}_{1/n}} \min_{\widehat{S}(\Delta \mathfrak{A}_n) \in \mathcal{Z}} \frac{\|\Delta \mathfrak{A}_n\|}{\|\dot{\Delta}_n\|} = \frac{\|\Delta \mathfrak{A}'_n\|}{\|\dot{\Delta}_n\|} \quad \text{and} \quad s_n = \min_{\Delta \mathfrak{A}_n \in \mathcal{G}_{1/n}} \min_{\widehat{S}(\Delta \mathfrak{A}_n) \in \mathcal{Z}} \frac{\|T_{\mathbf{p}} \dot{\Delta}_n\|}{\|\dot{\Delta}_n\|},$$

where n is sufficiently large, and then it follows from the first part of (36) that

$$c_n \geq s_n + \min_{\Delta \mathfrak{A}_n \in \mathcal{G}_{1/n}} \min_{\widehat{S}(\Delta \mathfrak{A}_n) \in \mathcal{Z}} \varepsilon_n \quad \text{and} \quad s_n \geq c_n + \min_{\Delta \mathfrak{A}_n \in \mathcal{G}_{1/n}} \min_{\widehat{S}(\Delta \mathfrak{A}_n) \in \mathcal{Z}} (-\varepsilon_n);$$

note that I used the subscript n for indicating the dependence of $\dot{\Delta}$ on $\Delta \mathfrak{A}_n$. From the second part of (36), it follows that we can bound

$$(39) \quad s_n - \vartheta_n \leq c_n \leq s_n + \vartheta_n, \quad \text{so} \quad (s_n - \vartheta_n)^{-1} \geq c_n^{-1} \geq (s_n + \vartheta_n)^{-1},$$

for sufficiently large n so that $s_n > \vartheta_n$. Note that $\vartheta_n \rightarrow 0$ as $n \rightarrow \infty$ because of (36), (31) and Corollary 6. Since $\dot{\Delta}_n \in \mathcal{K}^\perp$, it follows immediately from the Courant–Fisher minimax characterization of the least singular values [49, Theorem 4.2.11] that

$$0 < \varsigma_N(T_{\mathbf{p}}) \leq \frac{\|T_{\mathbf{p}} \dot{\Delta}_n\|}{\|\dot{\Delta}_n\|} \leq \varsigma_1(T_{\mathbf{p}}),$$

where $N = r(\Sigma) + 1$, $\varsigma_i(T_{\mathbf{p}})$ denotes the i th singular value of $T_{\mathbf{p}}$, and the first inequality is due to the assumption on the rank of $T_{\mathbf{p}}$. I claim that the sequence of s_n converges:

$$(40) \quad \lim_{n \rightarrow \infty} s_n = \lim_{\epsilon \rightarrow 0} \min_{\Delta \mathfrak{A} \in \mathcal{G}_\epsilon} \min_{\widehat{S}(\Delta \mathfrak{A}) \in \mathcal{Z}} \frac{\|T_{\mathbf{p}} \dot{\Delta}(\Delta \mathfrak{A}, \widehat{S}(\Delta \mathfrak{A}))\|}{\|\dot{\Delta}(\Delta \mathfrak{A}, \widehat{S}(\Delta \mathfrak{A}))\|} = \varsigma_N(T_{\mathbf{p}}).$$

This statement would conclude the proof because then all of the following limits are well-defined:

$$\lim_{n \rightarrow \infty} (s_n - \vartheta_n)^{-1} = (\varsigma_N(T_{\mathbf{p}}))^{-1} \quad \text{and} \quad \lim_{n \rightarrow \infty} (s_n + \vartheta_n)^{-1} = (\varsigma_N(T_{\mathbf{p}}))^{-1},$$

sandwiching c_n^{-1} in (39), so that by considering (38) we find

$$\lim_{n \rightarrow \infty} c_n^{-1} \leq \lim_{n \rightarrow \infty} \max_{\Delta \mathfrak{A}_n \in \mathcal{G}_{1/n}} \frac{\|\widetilde{\Delta \mathbf{p}}_n\|}{\|\Delta \mathfrak{A}_n\|} \leq \lim_{n \rightarrow \infty} c_n^{-1} \cdot \lim_{n \rightarrow \infty} \max_{\Delta \mathfrak{A}_n \in \mathcal{G}_{1/n}} (1 + \delta_n),$$

and as $\delta_n \rightarrow 0$, it follows that

$$\kappa_A = \lim_{\epsilon \rightarrow 0} \max_{\Delta \mathfrak{A} \in \mathcal{G}_\epsilon} \frac{d(\mathbf{p}, f^\dagger(\mathfrak{A} + \Delta \mathfrak{A}))}{\|\Delta \mathfrak{A}\|} = \lim_{\epsilon \rightarrow 0} \max_{\Delta \mathfrak{A} \in \mathcal{G}_\epsilon} \frac{\|\widetilde{\Delta \mathbf{p}}_n\|}{\|\Delta \mathfrak{A}_n\|} = (\varsigma_N(T_{\mathbf{p}}))^{-1},$$

which would conclude the proof.

Part III: The limit. For proving (40), we consider the specific perturbation $\Delta \mathbf{p}'_n = n^{-1} \mathbf{w}$, where \mathbf{w} is the right singular vector of $T_{\mathbf{p}}$ corresponding to the singular value $\varsigma_N(T_{\mathbf{p}})$. The perturbed tensor is $\mathfrak{A}_n = f(\mathbf{p} + \Delta \mathbf{p}'_n) = f(\mathbf{p} + n^{-1} \mathbf{w}) = \mathfrak{A} + \Delta \mathfrak{A}_n$. If n_0 is sufficiently large, then $\mathfrak{A}_n \subset \mathcal{N}$ for all $n > n_0$. Clearly, $\|\Delta \mathfrak{A}_n\| \rightarrow 0$ as $n \rightarrow \infty$. From the Taylor series expansion at the start of the proof, it is also clear that for large n , $\Delta \mathfrak{A}_n \neq 0$, as the norms of the higher-order terms will be dominated by $n^{-1} = \|n^{-1} \mathbf{w}\|$. Let $\Delta \mathbf{p}_n$ be the specific normalization of $\mathbf{p} + \Delta \mathbf{p}_n \in f^\dagger(\mathfrak{A} + \Delta \mathfrak{A}_n)$ that satisfies the properties imposed in (29). Let $\widehat{S}_n \in \arg \min_{S \in \mathcal{T}} \|\mathbf{p} - S(\mathbf{p} + \Delta \mathbf{p}_n)\|$ be arbitrary. Then, $\mathbf{p} + \widehat{\Delta \mathbf{p}}_n = \widehat{S}_n(\mathbf{p} + \Delta \mathbf{p}_n)$. By definition of $\Delta \mathbf{p}_n$ there exists some $T_n \in \mathcal{T}$ such that $\mathbf{p} + \Delta \mathbf{p}_n = T_n(\mathbf{p} + \Delta \mathbf{p}'_n)$. Hence, $\mathbf{p} + \widehat{\Delta \mathbf{p}}_n = \widehat{S}_n T_n(\mathbf{p} + n^{-1} \mathbf{w})$, so that

$$(41) \quad \begin{aligned} \|\widehat{\Delta \mathbf{p}}_n\| &= \min_{S \in \mathcal{T}} \|\mathbf{p} - S(\mathbf{p} + \Delta \mathbf{p}_n)\| = \min_{S \in \mathcal{T}} \|\mathbf{p} - S T_n(\mathbf{p} + n^{-1} \mathbf{w})\| \\ &= \min_{S \in \mathcal{T}} \|\mathbf{p} - S(\mathbf{p} + n^{-1} \mathbf{w})\| \leq \|\mathbf{p} - I(\mathbf{p} + n^{-1} \mathbf{w})\| = n^{-1}; \end{aligned}$$

the third equality is because \mathcal{T} is a group.

Let n be sufficiently large. Then, I claim that $\widehat{S}_n T_n = I + E_n$ where E_n is a diagonal matrix whose diagonal entries are bounded by $\mathcal{O}(n^{-1})$. Indeed, the proof of Proposition 4 shows that (41) may only be true if for all $k = 2, \dots, d$ and all $i = 1, \dots, r$ simultaneously we have that

$$(42) \quad \|\mathbf{a}_i^k - \theta_{k, \pi_i}(\mathbf{a}_{\pi_i}^k + n^{-1} \mathbf{w}_{\pi_i}^k)\| \leq n^{-1}, \quad \theta_{k, \pi_i} \in \mathbb{F} \setminus \{0\},$$

where $\mathbf{w} = \text{vecr}(\omega_1, \dots, \omega_r)$, $\omega_i = (\mathbf{w}_i^1, \dots, \mathbf{w}_i^d)$, and π is a permutation of $\{1, 2, \dots, r\}$. Since \mathbf{a}_i^k are constant vectors, it follows that this inequality can only be satisfied if

$$\mathbf{a}_i^k = \theta_{k, \pi_i} \mathbf{a}_{\pi_i}^k \quad \text{for all } k = 2, 3, \dots, d \text{ and } i = 1, 2, \dots, r.$$

Suppose that π is not the identity permutation, then there exists a $j \neq \pi_j$ for which it holds that

$$\begin{aligned} \mathbf{a}_j^1 \otimes \mathbf{a}_j^2 \otimes \dots \otimes \mathbf{a}_j^d + \mathbf{a}_{\pi_j}^1 \otimes \mathbf{a}_{\pi_j}^2 \otimes \dots \otimes \mathbf{a}_{\pi_j}^d &= \mathbf{a}_j^1 \otimes \mathbf{a}_j^2 \otimes \dots \otimes \mathbf{a}_j^d + \mathbf{a}_{\pi_j}^1 \otimes (\theta_{2, \pi_j}^{-1} \mathbf{a}_j^2) \otimes \dots \otimes (\theta_{d, \pi_j}^{-1} \mathbf{a}_j^d) \\ &= (\mathbf{a}_j^1 + \theta_{2, \pi_j}^{-1} \dots \theta_{d, \pi_j}^{-1} \mathbf{a}_{\pi_j}^1) \otimes \mathbf{a}_j^2 \otimes \dots \otimes \mathbf{a}_j^d, \end{aligned}$$

showing that \mathfrak{A} really has a decomposition of length at most $r - 1$, which is a contradiction. So π is the identity; hence (42) may be simplified to

$$(43) \quad \|(1 - \theta_{k, i}) \mathbf{a}_i^k - n^{-1} \theta_{k, i} \mathbf{w}_i^k\|^2 \leq n^{-2},$$

which should hold for all $i = 1, 2, \dots, r$ and $k = 2, 3, \dots, d$. It can be satisfied only if

$$(44) \quad 1 - 3n^{-1} \|\mathbf{a}_i^k\|^{-1} \leq \theta_{k, i} \leq 1 + 3n^{-1} \|\mathbf{a}_i^k\|^{-1}.$$

Note that (43) is a quadratic equation in $\theta_{k, i}$ with positive coefficient for $\theta_{k, i}^2$. So it suffices proving that inequality (43) is satisfied for $\theta_{k, i} = 1$ while it is no longer satisfied for $\theta_{k, i} = 1 \pm 3n^{-1} \|\mathbf{a}_i^k\|^{-1}$ to conclude that any $\theta_{k, i}$ satisfying (43) must be contained in the interval (44). For $\theta_{k, i} = 1$ one has $\|(1 - 1) \mathbf{a}_i^k - n^{-1} \mathbf{w}_i^k\| = n^{-1} \|\mathbf{w}_i^k\| \leq n^{-1}$, so (43) is satisfied. Pugging $\theta_{k, i} = 1 \pm 3n^{-1} \|\mathbf{a}_i^k\|^{-1}$ into (43), we get from the triangle inequality

$$\begin{aligned} \|(1 - \theta_{k, i}) \mathbf{a}_i^k - n^{-1} \theta_{k, i} \mathbf{w}_i^k\| &\geq 3n^{-1} \|\mathbf{a}_i^k\|^{-1} \|\mathbf{a}_i^k\| - n^{-1} |1 \pm 3n^{-1} \|\mathbf{a}_i^k\|^{-1}| \cdot \|\mathbf{w}_i^k\| \\ &\geq n^{-1} (3 - |1 \pm 3n^{-1} \|\mathbf{a}_i^k\|^{-1}|) = n^{-1} (2 \mp 3n^{-1} \|\mathbf{a}_i^k\|^{-1}) > n^{-1} \end{aligned}$$

where the last equality and inequality are valid only for large n , namely $n > 3 \|\mathbf{a}_i^k\|^{-1}$. Hence, $|\theta_{k, i} - 1| \leq 3n^{-1} \|\mathbf{a}_i^k\|^{-1}$ for $k = 2, \dots, d$. For $\theta_{1, i} = (\theta_{2, i} \dots \theta_{d, i})^{-1}$ we get the usual bound

$$\begin{aligned} |\theta_{1, i} - 1| &= \left| -1 + \prod_{k=2}^d (1 \pm 3n^{-1} \|\mathbf{a}_i^k\|^{-1})^{-1} \right| \leq -1 + 1 + \sum_{\kappa=1}^{\infty} \sum_{\|\ell\|_1=\kappa} \prod_{k=2}^d \left| 3n^{-1} \|\mathbf{a}_i^k\|^{-1} \right|^{\ell_k} \\ &\leq \sum_{\kappa=1}^{\infty} (3n^{-1} \chi_i)^\kappa \|\mathbf{1} \otimes \dots \otimes \mathbf{1}\|_1 \leq \sum_{\kappa=1}^{\infty} (3(d-1)n^{-1} \chi_i)^\kappa \leq 6(d-1)n^{-1} \chi_i, \end{aligned}$$

where in the third step the same argument surrounding the derivation of (20) was employed, and where in the last step we assumed $n \geq 6(d-1)\chi_i$. Letting $\mu_{k,i}$ be defined as

$$(45a) \quad \theta_{k,i} = 1 - \mu_{k,i}, \quad k = 2, 3, \dots, d, \quad i = 1, 2, \dots, r$$

$$(45b) \quad \theta_{1,i} = \theta_{2,i}^{-1} \cdots \theta_{d,i}^{-1} = 1 - \mu_{1,i} \quad i = 1, 2, \dots, r$$

we may write E_n explicitly as

$$E_n = \text{diag}(E_{n,1}, \dots, E_{n,r}) \quad \text{with } E_{n_i} = \text{diag}(-\mu_{1,i}I_{n_1}, -\mu_{2,i}I_{n_2}, \dots, -\mu_{d,i}I_{n_d}).$$

By the foregoing derivation $|\mu_{k,i}| = \mathcal{O}(n^{-1})$.

We now consider the following factorization

$$(46) \quad \mathbf{p} + \widetilde{\Delta \mathbf{p}}_n = \widehat{S}_n(\mathbf{p} + \Delta \mathbf{p}_n) = \widehat{S}_n T_n(\mathbf{p} + \Delta \mathbf{p}'_n) = (I + E_n)(\mathbf{p} + \Delta \mathbf{p}'_n) \\ = \mathbf{p} + \Delta \mathbf{p}'_n + E_n \mathbf{p} + E_n \Delta \mathbf{p}'_n \\ = \mathbf{p} + (\nabla_{n,1} + \nabla_{n,2}) + (\Delta_{n,1} + \Delta_{n,2} + \Delta \mathbf{p}'_n)$$

where $E_n \mathbf{p} = \nabla_{n,1} + \Delta_{n,1}$ and $E_n \Delta \mathbf{p}'_n = \nabla_{n,2} + \Delta_{n,2}$ with $\nabla_{n,i} \in \mathcal{K}$ and $\Delta_{n,i} \in \mathcal{K}^\perp$. One sees immediately that $\|\nabla_{n,1}\| \leq \|E_n\|_2 \|\mathbf{p}\| = \mathcal{O}(n^{-1})$, $\|\nabla_{n,2}\| \leq \|E_n\|_2 \|\Delta \mathbf{p}'_n\| = \mathcal{O}(n^{-2})$, and $\|\Delta_{n,2}\| \leq \|E_n\|_2 \|\Delta \mathbf{p}'_n\| = \mathcal{O}(n^{-2})$. The key difficulty lies in proving that $\|\Delta_{n,1}\| = \mathcal{O}(n^{-2})$. By definition,

$$\Delta_{n,1} = (I - K K^\dagger) E_n \mathbf{p} = (I - K \text{diag}(K_1^H K_1, \dots, K_r^H K_r)^{-1} K^H) E_n \mathbf{p}$$

so it suffices to prove that $\|(I - K_i K_i^\dagger) \mathbf{q}_i\| = \mathcal{O}(n^{-2})$, where

$$\mathbf{q}_i = \mathbf{q}_i^{(1)} = \begin{bmatrix} -\mu_{1,i} \mathbf{a}_i^1 \\ -\mu_{2,i} \mathbf{a}_i^2 \\ \vdots \\ -\mu_{d,i} \mathbf{a}_i^d \end{bmatrix}$$

with $\mu_{k,i}$ as in (45). By definition and some Maclaurin series, we know that

$$(47) \quad -\mu_{1,i} = -1 + \prod_{k=2}^d (1 - \mu_{k,i})^{-1} = \sum_{k=2}^d \mu_{k,i} + \sum_{\kappa=2} \sum_{\|\ell\|_1=\kappa} \prod_{k=2}^d \mu_{k,i}^{\ell_k} = \sum_{k=2}^d \mu_{k,i} + M_i;$$

it can be shown in the usual way that $|M_i| = \mathcal{O}(n^{-2})$. Let $\alpha_{k,i} = \|\mathbf{a}_i^k\|$. Then, we find $\mathbf{q}_i^{(2)} = K_i^H \mathbf{q}_i^{(1)}$.

The entries of $\mathbf{q}_i^{(2)}$ are

$$q_{k-1,i}^{(2)} = -\mu_{1,i} \alpha_{1,i}^2 + \mu_{k,i} \alpha_{k,i}^2, \quad k = 2, \dots, d.$$

For the next step $\mathbf{q}_i^{(3)} = (K_i^H K_i)^{-1} \mathbf{q}_i^{(2)}$, we recall from the Sherman–Morrison formula that

$$(K_i^H K_i)^{-1} = (\text{diag}(\alpha_{2,i}^2, \dots, \alpha_{d,i}^2) + \alpha_{1,i}^2 \mathbf{1} \mathbf{1}^T)^{-1} = \text{diag}(\alpha_{2,i}^{-2}, \dots, \alpha_{d,i}^{-2}) - \frac{\alpha_{1,i}^2}{1 + \alpha_{1,i}^2 \sum_{j=2}^d \alpha_{j,i}^{-2}} \mathbf{h} \mathbf{h}^T,$$

where $\mathbf{h}_i^T = [\alpha_{2,i}^{-2} \cdots \alpha_{d,i}^{-2}]$. Let $\widehat{\alpha}_i = \alpha_{1,i}^2 \sum_{j=2}^d \alpha_{j,i}^{-2}$. It is then straightforward to establish that

$$\mathbf{h}_i^T \mathbf{q}_i^{(2)} = -\mu_{1,i} \widehat{\alpha}_i + \sum_{k=2}^d \mu_{k,i},$$

so that the elements of $\mathbf{q}_i^{(3)}$ are given by

$$q_{k-1,i}^{(3)} = -\mu_{1,i} \alpha_{1,i}^2 \alpha_{k,i}^{-2} + \mu_{k,i} - \frac{\alpha_{1,i}^2 \alpha_{k,i}^{-2} (-\mu_{1,i} \widehat{\alpha}_i + \sum_{k=2}^d \mu_{k,i})}{1 + \widehat{\alpha}_i} \\ = -\mu_{1,i} \alpha_{1,i}^2 \alpha_{k,i}^{-2} + \mu_{k,i} - \frac{\alpha_{1,i}^2 \alpha_{k,i}^{-2} (1 + \widehat{\alpha}_i) \sum_{k=2}^d \mu_{k,i} + \alpha_{1,i}^2 \alpha_{k,i}^{-2} \widehat{\alpha}_i M_i}{1 + \widehat{\alpha}_i} \\ = \mu_{k,i} + \alpha_{1,i}^2 \alpha_{k,i}^{-2} M_i - \alpha_{1,i}^2 \alpha_{k,i}^{-2} \widehat{\alpha}_i (1 + \widehat{\alpha}_i)^{-1} M_i \\ = \mu_{k,i} + M_i \alpha_{1,i}^2 \alpha_{k,i}^{-2} (1 - \widehat{\alpha}_i (1 + \widehat{\alpha}_i)^{-1}) = \mu_{k,i} + M_{i,k},$$

where in the second and third step (47) was used. Finally $\mathbf{q}_i^{(4)} = \mathbf{q}_i^{(1)} - K_i \mathbf{q}_i^{(3)} = (I - K_i K_i^\dagger) \mathbf{q}_i$ where

$$\mathbf{q}_i^{(4)} = \begin{bmatrix} ((-\mu_{1,i} - \sum_{k=2}^d (\mu_{k,i} + M_{i,k})) \mathbf{a}_i^1) \\ (-\mu_{2,i} + \mu_{2,i} + M_{i,2}) \mathbf{a}_i^2 \\ \vdots \\ (-\mu_{d,i} + \mu_{d,i} + M_{i,d}) \mathbf{a}_i^d \end{bmatrix} = \begin{bmatrix} (M_i - \sum_{k=2}^d M_{i,k}) \mathbf{a}_i^1 \\ M_{i,2} \mathbf{a}_i^2 \\ \vdots \\ M_{i,d} \mathbf{a}_i^d \end{bmatrix},$$

from which it is obvious that $\|\Delta_{n,1}\| = \mathcal{O}(n^{-2})$.

Continuing from (46), we can apply the Iterated Scaling Lemma to $\mathbf{p} + \nabla_n$, where $\nabla_n = \nabla_{n,1} + \nabla_{n,2}$ and $\|\nabla_n\| = \mathcal{O}(n^{-1})$, thusly obtaining

$$\mathbf{p} + \widetilde{\Delta \mathbf{p}}_n = \dot{\mathbf{p}}_n + \Delta'_n + \Delta_{n,1} + \Delta_{n,2} + \Delta \mathbf{p}'_n = \dot{\mathbf{p}}_n + n^{-1} \mathbf{w} + n^{-2} \mathbf{x}_n,$$

where $\Delta'_n \in \mathcal{K}^\perp$ is of norm $\mathcal{O}(n^{-2})$ for sufficiently large n , and \mathbf{x}_n is a vector whose norm can be uniformly bounded by a constant C . By definition, $\dot{\Delta}_n = n^{-1} \mathbf{w} + n^{-2} \mathbf{x}_n$ and $T_{\mathbf{p}} \mathbf{w} = \varsigma_N \mathbf{w}$. From the sandwiching

$$\frac{\varsigma_N - n^{-1} \|T_{\mathbf{p}}\|_2 \|\mathbf{x}_n\|}{1 + n^{-1} \|\mathbf{x}_n\|} \leq \frac{n^{-1} \|T_{\mathbf{p}}(\mathbf{w} + n^{-1} \mathbf{x}_n)\|}{n^{-1} \|\mathbf{w} + n^{-1} \mathbf{x}_n\|} \leq \frac{\varsigma_N + n^{-1} \|T_{\mathbf{p}}\|_2 \|\mathbf{x}_n\|}{1 - n^{-1} \|\mathbf{x}_n\|},$$

it follows that $\|T_{\mathbf{p}} \dot{\Delta}(\Delta \mathfrak{A}_n, \bullet)\| / \|\dot{\Delta}(\Delta \mathfrak{A}_n, \bullet)\|$, where “ \bullet ” indicates an arbitrary $\widehat{S}(\Delta \mathfrak{A}_n) \in \mathcal{Z}$, tends to $\varsigma_N(T_{\mathbf{p}})$ as $n \rightarrow \infty$, concluding the proof. \square

A relative condition number may then be defined as follows.

Definition 3 (Condition number). Let \mathbf{p} and \mathfrak{A} be as in Theorem 7. Then, the relative condition number of the tensor rank decomposition problem $\mathfrak{A} = f(\mathbf{p})$ at \mathbf{p} is

$$\kappa = \kappa_A \cdot \frac{\|\mathfrak{A}\|}{\|\mathbf{p}\|}.$$

7. THE NORM-BALANCED CONDITION NUMBER

It should be remarked that the absolute and relative condition numbers considered in the previous section are defined at \mathbf{p} rather than at \mathfrak{A} . The reason is that their value depends on the particular representative $\mathbf{p} \in f^\dagger(\mathfrak{A})$ that was chosen. With the foregoing definition of an absolute and a relative condition number, we could now define the condition number of the entire fiber $f^\dagger(\mathfrak{A})$ in several ways, each with its particular interpretation. For instance, one could define the condition number of f^\dagger at \mathfrak{A} to be the maximum, minimum or (weighted) average condition number of all $\mathbf{p} \in f^\dagger$, which would respectively measure the conditioning of the worst possible representative, the best possible representative, and the (weighted) average representative. I believe that all of the foregoing definitions are very interesting from a mathematical viewpoint, however I am skeptical that such definitions correspond naturally with the current practice in data-analytic applications involving the tensor rank decomposition. Through private conversation with G. Tomasi [74], it appeared to me that in many applications, especially originating from physical measurements, there is no a priori reason to prefer one scaling of the factors over another. Therefore, I believe that a natural choice for the condition number of f^\dagger at \mathfrak{A} is the condition number of the norm-balanced case, which is defined next.

Definition 4. (Norm-balanced condition number) Let $\alpha_i \in \mathbb{R}_+$ be strictly positive,¹¹ and let

$$\mathfrak{A} = \sum_{i=1}^r \alpha_i \mathbf{u}_i^1 \otimes \cdots \otimes \mathbf{u}_i^d \quad \text{with } \|\mathbf{u}_i^k\| = 1$$

be a robustly r -identifiable tensor rank decomposition in $\mathbb{F}^{n_1} \otimes \cdots \otimes \mathbb{F}^{n_d}$. Then, its condition number is defined to be the condition number of the norm-balanced representatives

$$p_i = (\alpha_i^{1/d} \mathbf{u}_i^1, \alpha_i^{1/d} \mathbf{u}_i^2, \dots, \alpha_i^{1/d} \mathbf{u}_i^d), \quad i = 1, 2, \dots, r.$$

¹¹Every rank- r tensor can be represented as such, by absorbing the sign into, e.g., \mathbf{u}_i^1 .

That is, the norm-balanced absolute and relative condition numbers of f^\dagger at \mathfrak{A} are respectively

$$\kappa_A = \varsigma_N^{-1} \quad \text{and} \quad \kappa = \varsigma_N^{-1} \frac{\|\mathfrak{A}\|}{\|\mathbf{p}\|},$$

where $\mathbf{p} = \text{vecr}(p_1, p_2, \dots, p_r)$ and ς_N is the $N = r(\Sigma + 1)$ th largest singular value of $T_{\mathbf{p}}$.

Whenever I refer to respectively *the* absolute and relative condition number of a tensor, I mean the norm-balanced absolute and relative condition number, respectively.

In the remainder of the paper, only the norm-balanced condition number of f^\dagger will be considered.

The norm-balanced condition number has another practical advantage over the other choices: it is very easy to compute when a decomposition of the tensor is given, requiring only the least singular value of one Terracini's matrix; for completeness, such a procedure is presented as Algorithm 1. Lines 1 to 9 implement the norm-balancing of the input tensor rank decomposition. The required Terracini matrix is constructed in lines 10 to 13. Finally, the condition number is computed in line 15. A concrete implementation of this algorithm in Matlab/Octave is included in Appendix A.

Algorithm 1: Computing the norm-balanced relative condition number

input : A robustly r -identifiable tensor $\mathfrak{A} = \sum_{i=1}^r \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^d \in \mathbb{F}^{n_1} \otimes \dots \otimes \mathbb{F}^{n_d}$
output: The norm-balanced condition number $\kappa \in \mathbb{R}_+$ of \mathfrak{A} .

```

1  $\zeta \leftarrow 1$ 
2 for  $i = 1, 2, \dots, r$  do
3    $\gamma \leftarrow 1$ 
4   for  $k = 1, 2, \dots, d$  do
5      $\eta \leftarrow \|\mathbf{a}_i^k\|$ ;  $\gamma \leftarrow \gamma \cdot \eta^{1/d}$ ;  $\mathbf{a}_i^k \leftarrow \mathbf{a}_i^k / \eta$ 
6   end
7    $\mathbf{a}_i^k \leftarrow \gamma \mathbf{a}_i^k$ ,  $k = 1, 2, \dots, d$ 
8    $\zeta \leftarrow d\gamma^2$ 
9 end
10  $T \leftarrow []$ 
11 for  $i = 1, 2, \dots, r$  do
12    $T \leftarrow [T \quad I_{n_1} \otimes \mathbf{a}_i^2 \otimes \dots \otimes \mathbf{a}_i^d \quad \dots \quad \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^{d-1} \otimes I_{n_d}]$ 
13 end
14  $N \leftarrow r(n_1 + \dots + n_d - d + 1)$ 
15  $\kappa \leftarrow (\varsigma_N(T)\sqrt{\zeta})^{-1} \cdot \|\mathfrak{A}\|$ 
```

The computational complexity of computing the norm-balanced condition number is determined as follows. Note that we may always assume that $r\Sigma < \Pi$, because otherwise the input cannot be robustly r -identifiable. Since the condition number is ∞ in this case no involved computations are necessary. It is straightforward to verify that the computation of the least singular value in line 15 of Algorithm 1 contributes the dominant factor to the time complexity analysis. For a general $m \times n$ matrix with $n \leq m$, computing the least singular value costs $\mathcal{O}(mn^2)$ operations. The asymptotic space complexity is determined by the cost of storing Terracini's matrix, which is of size $\Pi \times r(\Sigma + d)$. This proves the following elementary result.

Proposition 8. *If a length r decomposition of a tensor in $\mathbb{F}^{n_1} \otimes \dots \otimes \mathbb{F}^{n_d}$ is supplied to Algorithm 1 as input, then its time and space complexity are respectively*

$$\mathcal{O}(\Pi(r(\Sigma + d))^2) = \mathcal{O}(\Pi^3) \quad \text{and} \quad \mathcal{O}(r(\Sigma + d)\Pi) = \mathcal{O}(\Pi^2),$$

where $\Pi = \prod_{i=1}^d n_i$ and $\Sigma + d = \sum_{i=1}^d n_i$.

8. BASIC PROPERTIES OF THE NORM-BALANCED CONDITION NUMBER

We proceed by investigating several basic properties of the norm-balanced condition number, starting with three desirable properties that are also exhibited by the familiar condition number of a matrix, namely its continuity (under mild conditions), and its invariance under scaling and orthogonal change of bases.

Proposition 9. *The absolute and relative condition numbers are continuous in a neighborhood of a robustly r -identifiable tensor.*

Proof. This follows immediately from the continuity of Terracini's matrix $T_{\mathbf{p}}$ in the parameters \mathbf{p} , which is a consequence from the assumption of regularity in the definition of robust r -identifiability. \square

Proposition 10. *The relative condition number is scale-invariant.*

Proof. The claim is that $\kappa(\mathfrak{A}) = \kappa(\beta\mathfrak{A})$ for all $\beta \in \mathbb{F} \setminus \{0\}$. Note that $\beta\mathfrak{A} = \alpha^d\mathfrak{A} = f(\alpha\mathbf{p})$, so $\|\alpha^d\mathfrak{A}\|/\|\alpha\mathbf{p}\| = |\alpha|^{d-1}\|\mathfrak{A}\|/\|\mathbf{p}\|$. Terracini's matrix corresponding to the points represented by $\mathbf{q} = \alpha\mathbf{p}$ is verified to be $T_{\mathbf{q}} = \alpha^{d-1}T_{\mathbf{p}}$, where $T_{\mathbf{p}}$ is Terracini's matrix corresponding to \mathbf{p} . Hence, $\varsigma_N(T_{\mathbf{q}})^{-1} = (|\alpha|^{d-1}\varsigma_N(T_{\mathbf{p}}))^{-1} = |\alpha|^{1-d}\kappa_A(T_{\mathbf{p}})$. This concludes the proof for $\mathbb{F} = \mathbb{C}$. In the real case, whenever d is even, one should exploit for $\beta < 0$ the equality $\beta\mathfrak{A} = -\alpha^d\mathfrak{A} = f(\alpha\mathbf{p}')$, where $\mathbf{p}' = \text{vecr}(p'_1, \dots, p'_r)$ with $p'_i = (-\mathbf{a}_i^1, \mathbf{a}_i^2, \dots, \mathbf{a}_i^d)$. Then, Terracini's matrix corresponding to $\mathbf{q}' = \alpha\mathbf{p}'$ is $T_{\mathbf{q}'} = \alpha^{d-1}T_{\mathbf{p}}S$, where S is a diagonal matrix whose diagonal entries are ± 1 . Since S is an orthogonal matrix, $\varsigma_N(T_{\mathbf{p}}) = \varsigma_N(T_{\mathbf{p}}S)$ and thusly is the proof concluded. \square

Proposition 11. *The absolute and relative condition number are orthogonally invariant.*

Proof. Let \mathbf{p} and \mathfrak{A} be as in Theorem 7. The claim is that κ_A and κ are invariants of the $(O(n_1) \times \dots \times O(n_d))$ -orbit of \mathfrak{A} . In other words, considering $\mathfrak{A} \in \mathbb{F}^\Pi$, the statement is that

$$\kappa(\mathfrak{A}) = \kappa((Q_1 \otimes \dots \otimes Q_d)\mathfrak{A})$$

where $Q_i \in O(n_i)$ are orthogonal matrices in $\mathbb{F}^{n_i \times n_i}$ with respect to the Hermitian inner product. As $Q = Q_1 \otimes \dots \otimes Q_d$ is orthogonal, $\|\mathfrak{A}\| = \|Q\mathfrak{A}\|$. Since

$$Q\mathfrak{A} = (Q_1 \otimes \dots \otimes Q_d) \sum_{i=1}^r \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^d = \sum_{i=1}^r Q_1 \mathbf{a}_i^1 \otimes \dots \otimes Q_d \mathbf{a}_i^d,$$

it follows that $Q\mathfrak{A} = f(\text{diag}(Q_1, \dots, Q_d, \dots, Q_1, \dots, Q_d)\mathbf{p}) = f(U\mathbf{p})$. As U is orthogonal, $\|\mathbf{p}\| = \|U\mathbf{p}\|$. Finally, Terracini's matrix corresponding to $\mathbf{p}' = U\mathbf{p}$ is given explicitly by

$$T_{\mathbf{p}'} = [T'_1 \quad \dots \quad T'_r], \quad \text{where } T'_i = [I \otimes Q_2 \mathbf{a}_i^2 \otimes \dots \otimes Q_d \mathbf{a}_i^d \quad \dots \quad Q_1 \mathbf{a}_i^1 \otimes \dots \otimes Q_{d-1} \mathbf{a}_i^{d-1} \otimes I].$$

Multiplying T'_i on the right by $D = \text{diag}(Q_1, Q_2, \dots, Q_d)$ results in $T'_i D = Q T_i$, so that $Q T_{\mathbf{p}} = T_{\mathbf{p}'} \text{diag}(D, \dots, D)$. Since Q and $\text{diag}(D, \dots, D)$ are orthogonal, one obtains

$$\varsigma_N(T_{\mathbf{p}}) = \varsigma_N(Q T_{\mathbf{p}}) = \varsigma_N(T_{\mathbf{p}'} \text{diag}(D, \dots, D)) = \varsigma_N(T_{\mathbf{p}'}),$$

concluding the proof. \square

The condition number of a rank-1 tensor has an explicit expression; in fact, the relative condition number depends only on the order d of the tensor, and *decreases* as d increases. As rank-1 tensors are robustly r -identifiable for all $d \geq 2$, the result even applies to matrices.

Proposition 12. *A rank-1 tensor $\mathfrak{A} = \alpha \mathbf{a}^1 \otimes \dots \otimes \mathbf{a}^d \in \mathbb{F}^{n_1} \otimes \dots \otimes \mathbb{F}^{n_d}$ with $\|\mathbf{a}^k\| = 1$ and $\alpha \in \mathbb{R}_+$ strictly positive has absolute condition number $\kappa_A(\mathfrak{A}) = \alpha^{1/d-1}$ and relative condition number $\kappa(\mathfrak{A}) = \sqrt{d-1}$.*

Proof. Let $\mathbf{p}' = \text{vec}(\alpha^{1/d} \mathbf{a}^1, \alpha^{1/d} \mathbf{a}^2, \dots, \alpha^{1/d} \mathbf{a}^d)$ and $\mathbf{p} = \text{vec}(\mathbf{a}^1, \mathbf{a}^2, \dots, \mathbf{a}^d)$. Since the decomposition is norm-balanced, one verifies that $T_{\mathbf{p}'} = \alpha^{1-1/d} T_{\mathbf{p}}$. Hence, finding the singular values of $T_{\mathbf{p}}$ suffices. The largest singular value of $T_{\mathbf{p}}$ is given by

$$\varsigma_1(T_{\mathbf{p}}) = \max_{\|\mathbf{v}\|=1} \|T_{\mathbf{p}} \mathbf{v}\| = \max_{\|\mathbf{v}\|=1} \|\mathbf{v}^1 \otimes \mathbf{a}^2 \otimes \dots \otimes \mathbf{a}^d + \dots + \mathbf{a}^1 \otimes \dots \otimes \mathbf{a}^{d-1} \otimes \mathbf{v}^d\|$$

$$\begin{aligned}
&\leq \max_{\|\mathbf{v}\|=1} (\|\mathbf{v}^1 \otimes \mathbf{a}^2 \otimes \cdots \otimes \mathbf{a}^d\| + \cdots + \|\mathbf{a}^1 \otimes \cdots \otimes \mathbf{a}^{d-1} \otimes \mathbf{v}^d\|) \\
&= \max_{\|\mathbf{v}\|=1} (\|\mathbf{v}^1\| + \cdots + \|\mathbf{v}^d\|) = \max_{\|\mathbf{v}^1\|^2 + \cdots + \|\mathbf{v}^d\|^2 = 1} (\|\mathbf{v}^1\| + \cdots + \|\mathbf{v}^d\|) \\
&= \max_{\nu_1^2 + \cdots + \nu_d^2 = 1} (|\nu_1| + \cdots + |\nu_d|) = \max_{\|\boldsymbol{\nu}\|=1} \|\boldsymbol{\nu}\|_1 \leq \sqrt{d} \|\boldsymbol{\nu}\|_2 = \sqrt{d},
\end{aligned}$$

where $\mathbf{v}^T = [(\mathbf{v}^1)^T \cdots (\mathbf{v}^d)^T]$ and $\boldsymbol{\nu} = [\nu_1 \cdots \nu_d]$. If we take $\mathbf{v}_1 = \sqrt{d^{-1}} \text{vec}(\mathbf{a}^1, \dots, \mathbf{a}^d)$, then it follows that $\|T_{\mathbf{p}} \mathbf{v}_1\| = \|d\sqrt{d^{-1}} \mathbf{a}^1 \otimes \cdots \otimes \mathbf{a}^d\| = \sqrt{d}$. That is, \mathbf{v}_1 is a right singular vector corresponding to the largest singular value $\varsigma_1(T_{\mathbf{p}}) = \sqrt{d}$. Since $T_{\mathbf{p}} \mathbf{v}_1 = \sqrt{d} \mathbf{a}^1 \otimes \cdots \otimes \mathbf{a}^d = \sqrt{d} \mathbf{u}_1$, it follows that \mathbf{u}_1 is a left singular vector corresponding to $\varsigma_1(T_{\mathbf{p}})$. Deflating the largest singular tuple from $T_{\mathbf{p}}$, we find

$$\widehat{T}_{\mathbf{p}} = T_{\mathbf{p}} - \sqrt{d} \mathbf{u}_1 \mathbf{v}_1^H = [(I - \mathbf{a}^1(\mathbf{a}^1)^H) \otimes \mathbf{a}^2 \otimes \cdots \otimes \mathbf{a}^d \quad \cdots \quad \mathbf{a}^1 \otimes \cdots \otimes \mathbf{a}^{d-1} \otimes (I - \mathbf{a}^d(\mathbf{a}^d)^H)].$$

The singular values of the above matrix are the square roots of the eigenvalues of its Gram matrix, which has a particularly pleasing structure as it is block diagonal,

$$\widehat{T}_{\mathbf{p}}^H \widehat{T}_{\mathbf{p}} = \text{diag}((I - \mathbf{a}^1(\mathbf{a}^1)^H)^2, (I - \mathbf{a}^2(\mathbf{a}^2)^H)^2, \dots, (I - \mathbf{a}^d(\mathbf{a}^d)^H)^2),$$

which is a consequence from the fact that $I - \mathbf{a}^k(\mathbf{a}^k)^H$ is a projector onto the orthogonal complement of \mathbf{a}^k , so that applying it to the latter yields 0. Since $I - \mathbf{a}^k(\mathbf{a}^k)^H$ is a projector, it is idempotent, and so the eigenvalues of $\widehat{T}_{\mathbf{p}}^H \widehat{T}_{\mathbf{p}}$ are the union of the eigenvalues of the $I - \mathbf{a}^k(\mathbf{a}^k)^H$ for $k = 1, \dots, d$. The only possible eigenvalues of a projector are 1 and 0; so eigenvalue 1 has multiplicity equal to the rank of the projector, which is readily verified to be $n_k - 1$. Hence, the singular values of $\widehat{T}_{\mathbf{p}}$ are 1 with multiplicity Σ and 0 with multiplicity d . Then, the singular values of $T_{\mathbf{p}}$ are \sqrt{d} , 1 with multiplicity Σ , and 0 with multiplicity $d - 1$. We can conclude that the $N = (\Sigma + 1)$ th singular value of $T_{\mathbf{p}}$ is $\alpha^{1-1/d}$, so that the absolute condition number is $\kappa_A = \alpha^{1/d-1}$. Finally,

$$\kappa = \kappa_A \frac{\|\alpha \mathbf{a}^1 \otimes \cdots \otimes \mathbf{a}^d\|}{\|\mathbf{p}'\|} = \alpha^{1/d-1} \cdot \alpha \cdot (\sqrt{d} \alpha^{1/d})^{-1} = \sqrt{d^{-1}},$$

concluding the proof. \square

The following result analytically characterizes the singular values of Terracini's matrix corresponding to a norm-balanced rank-1 tensor; it is an immediate consequence from the foregoing proof.

Corollary 13. *Let $p = (\mathbf{a}^1, \mathbf{a}^2, \dots, \mathbf{a}^d)$ with $\|\mathbf{a}^1\| = \cdots = \|\mathbf{a}^d\| = \alpha^{1/d} > 0$, and let $\mathbf{p} = \text{vec}(p)$. Then, the singular values of Terracini's matrix are*

$$\varsigma(T_{\mathbf{p}}) = \{\sqrt{d} \alpha^{1-1/d}, \underbrace{\alpha^{1-1/d}, \dots, \alpha^{1-1/d}}_{\Sigma}, \underbrace{0, \dots, 0}_{d-1}\}.$$

Based on the foregoing characterization of the singular values of Terracini's matrix associated with one rank-1 tensor, we can determine explicitly the condition number of the class of *weak k -orthogonal tensors* with $k \geq 3$. Both strong and weak 2-orthogonality were considered in [78] from whence it name derives. A tensor rank decomposition is said to be weak k -orthogonal if for every pair of rank-1 tensors, there is orthogonality in k factors, however the factors in which orthogonality occurs need not be the same for different pairs. It was shown in [78] that a weak 2-orthogonal tensor rank decomposition is a necessary condition for the existence of a certain generalization of the Eckart–Young theorem for matrices, and a particular subclass of weak 2-orthogonal tensors, i.e., those with a strong 2-orthogonal decomposition, in fact admits a tensor-equivalent of the Eckart–Young theorem. In the next result, it is shown that tensors with a weak 3-orthogonal decomposition admit an easy expression for their condition number.

Proposition 14. *Let $\alpha_i \in \mathbb{R}_+$ be sorted as $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_r > 0$, and let*

$$\mathfrak{A} = \sum_{i=1}^r \alpha_i \mathbf{v}_i^1 \otimes \cdots \otimes \mathbf{v}_i^d \quad \text{with } \|\mathbf{v}_i^k\| = 1$$

be a robustly r -identifiable weak 3-orthogonal tensor:

$$\forall i < j : \exists 1 \leq k_1 < k_2 < k_3 \leq d : \langle \mathbf{v}_i^{k_1}, \mathbf{v}_j^{k_1} \rangle = \langle \mathbf{v}_i^{k_2}, \mathbf{v}_j^{k_2} \rangle = \langle \mathbf{v}_i^{k_3}, \mathbf{v}_j^{k_3} \rangle = 0,$$

where $\langle \cdot, \cdot \rangle$ is the Hermitian inner product. Then,

$$\kappa_A(\mathfrak{A}) = \alpha_r^{-1+1/d} \quad \text{and} \quad \kappa(\mathfrak{A}) = \alpha_r^{-1+1/d} \sqrt{\sum_{i=1}^r \alpha_i^2} / \sqrt{\sum_{i=1}^r d \alpha_i^{2/d}}.$$

Proof. Let $\mathbf{p} = \text{vecr}(p_1, \dots, p_r)$ with p_i the norm-balanced representatives as in Definition 4. By weak 3-orthogonality, we have for every $1 \leq i < j \leq r$ that there exist $k_1 < k_2$ so that

$$\begin{aligned} (\mathbf{v}_i^1 \otimes \dots \otimes \mathbf{v}_i^{k_1-1} \otimes I \otimes \mathbf{v}_i^{k_1+1} \otimes \dots \otimes \mathbf{v}_i^d)^H (\mathbf{v}_j^1 \otimes \dots \otimes \mathbf{v}_j^{k_2-1} \otimes I \otimes \mathbf{v}_j^{k_2+1} \otimes \dots \otimes \mathbf{v}_j^d) \\ = \mathbf{v}_j^{k_1} (\mathbf{v}_i^{k_2})^H \cdot \prod_{\substack{k=1, \\ k \neq k_1, k_2}}^d \langle \mathbf{v}_i^k, \mathbf{v}_j^k \rangle = 0, \end{aligned}$$

because the product excludes only 2 inner products so that there is always one inner product left with orthogonality. Then the foregoing entails that $T_i^H T_j = 0$ for all $i \neq j$, so that

$$T_{\mathbf{p}}^H T_{\mathbf{p}} = \text{diag}(T_1^H T_1, T_2^H T_2, \dots, T_r^H T_r),$$

where T_i is as in (4). The singular values of $T_{\mathbf{p}}$ thus coincide with the union of the set of singular values of each of the T_i 's. From Corollary 13 it follows that

$$\varsigma(T_{\mathbf{p}}) = \bigcup_{i=1}^r \{\sqrt{d} \alpha_i^{1-1/d}, \alpha_i^{1-1/d}, \dots, \alpha_i^{1-1/d}, 0, \dots, 0\};$$

the N th singular value is thus $\varsigma_N(T_{\mathbf{p}}) = \alpha_r^{1-1/d}$ by our assumption on the order of the α_i . The norm $\|\mathfrak{A}\|^2 = \sum_{i=1}^r \alpha_i^2$ because of the orthogonality of the $\mathbf{v}_i^1 \otimes \dots \otimes \mathbf{v}_i^d$'s. The norm of \mathbf{p} is clear. \square

Remark 15. The foregoing class of weak 3-orthogonal tensors includes rank-1 tensors and orthogonally decomposable tensors (odeco) of order $d \geq 3$ [14, 22, 51, 64, 78, 82].

As the relative condition number of order- d rank-1 tensors is $d^{-1/2}$, we can safely call such tensors *perfectly conditioned*: the relative error in the parameters of the decomposition is *less* than the relative error of the tensor. In fact, it is moderately decreasing with the order of the tensor. Recall that the relative condition number is defined in such a way that if $\kappa = 10^k$, then approximately k significant digits of accuracy are lost in the parameters of the decompositions $f^\dagger(\mathfrak{A})$ and $f^\dagger(\mathfrak{A} + \Delta\mathfrak{A})$ relative to the error between the tensors \mathfrak{A} and $\mathfrak{A} + \Delta\mathfrak{A}$ for sufficiently small $\|\Delta\mathfrak{A}\|$. Whenever $\log_{10}(\kappa)$ is small relative to $-\log_{10}(\|\Delta\mathfrak{A}\|/\|\mathfrak{A}\|)$, one may call \mathfrak{A} *well-conditioned*. If the former is large relative to the former, then \mathfrak{A} is called *ill-conditioned*. If $\kappa\|\Delta\mathfrak{A}\| > \|\mathfrak{A}\|$, then all significant digits may be lost in the output $f^\dagger(\mathfrak{A})$. The next result bounds the relative condition number of any robustly r -identifiable tensor from below.

Proposition 16. The relative condition number of robustly r -identifiable tensors in $\mathbb{F}^{n_1} \otimes \dots \otimes \mathbb{F}^{n_d}$ is bounded from below by d^{-1} .

Proof. Let \mathbf{p} and \mathfrak{A} be as in Theorem 7. Then,

$$\kappa(\mathfrak{A}) = \varsigma_N(T_{\mathbf{p}})^{-1} \frac{\|\mathfrak{A}\|}{\|\mathbf{p}\|} = \varsigma_N(T_{\mathbf{p}})^{-1} \frac{\|d^{-1} T_{\mathbf{p}} \mathbf{p}\|}{\|\mathbf{p}\|} = d^{-1} \varsigma_N(T_{\mathbf{p}})^{-1} \frac{\|T_{\mathbf{p}} \mathbf{p}\|}{\|\mathbf{p}\|}.$$

Let K be as in (5). Then by assumption on the rank of $T_{\mathbf{p}}$, $KK^\dagger = K(K^H K)^{-1}K^H$ is a projector onto the kernel of $T_{\mathbf{p}}$. Let $\mathbf{p}_i^T = [(\mathbf{a}_i^1)^T \dots (\mathbf{a}_i^d)^T]$, and then it follows that $K_i^H \mathbf{p}_i = 0$. Hence, $K^H \mathbf{p} = 0 = KK^\dagger \mathbf{p}$. So \mathbf{p} is contained in $\text{span}(K)^\perp$. Therefore, $\|T_{\mathbf{p}} \mathbf{p}\|/\|\mathbf{p}\| \geq \varsigma_N(T_{\mathbf{p}})$ and the result follows. \square

Remark 17. *It is unknown to me whether this lower bound may be sharp, because it is neither sharp for rank-1 tensors nor for weak 3-orthogonal tensors with all $\alpha_i = \alpha \neq 0$. I cannot presently think of tensors that would be even more well-behaved than completely orthogonal tensors all of whose terms are of equal norm; these tensors have condition number $\sqrt{d^{-1}}$.*

9. NUMERICAL EXAMPLES

In this section, some experiments are performed with the basic Matlab/Octave implementation of the proposed algorithm for computing the norm-balanced condition number that is included in the ancillary files accompanying this paper. All of the experiments were performed using Matlab R2015a on a computer system consisting of an Intel Core i7-5600U CPU, clocked at 2.6GHz, and 8GB of main memory. Some of the experiments employed Tensorlab v2.02 [68] for computing approximate tensor rank decompositions.

9.1. The Iterated Scaling Lemma. Let us start by experimentally investigating Lemma 5. It claims that a perturbation of the vectorized factor matrices \mathbf{p} in the direction of the kernel of Terracini's matrix $T_{\mathbf{p}}$ may be replaced with a new set of vectorized factor matrices representing the same tensor plus a much smaller perturbation perpendicular to the kernel of $T_{\mathbf{p}}$. For simplicity, take the following rank-2 tensor of size $3 \times 3 \times 2$, which is robustly r -identifiable by Kruskal's theorem:

$$\mathbf{A} = \begin{bmatrix} 2 & 0 \\ -1 & 1 \\ 0 & 2 \end{bmatrix}; \quad \mathbf{B} = \begin{bmatrix} -1 & -2 \\ 2 & 0 \\ 0 & 1 \end{bmatrix}; \quad \mathbf{C} = \begin{bmatrix} 1 & -2 \\ 2 & 1 \end{bmatrix};$$

These factor matrices are norm-balanced, and the corresponding Terracini's matrix can be constructed with the `cpdcond` algorithm that is included in the ancillary files:

$$[\mathbf{k}, \mathbf{F}, \mathbf{Tr}] = \text{cpdcond}(\{\mathbf{A}, \mathbf{B}, \mathbf{C}\});$$

The relative condition number is approximately $\kappa = 0.769$. Next, a random perturbation of unit norm in the direction of the kernel of \mathbf{Tr} is constructed as follows:

$$[\tilde{\cdot}, \tilde{\cdot}, \mathbf{V}] = \text{svd}(\mathbf{Tr}); \quad \mathbf{v} = \text{orth}(\mathbf{V}(:, 13:16) * \text{randn}(4, 1));$$

Now, the vectorized factor matrices are represented by the vector


$$\mathbf{p}^T = \begin{bmatrix} 2 & -1 & 0 & -1 & 2 & 0 & 1 & 2 & 0 & 1 & 2 & -2 & 0 & 1 & -2 & 1 \end{bmatrix}.$$

Applying the algorithm in Part I of the proof of the Iterated Scaling Lemma is accomplished as follows:

```
nabla = 1e-2*v; [q, delta] = isl(p, nabla, [3 3 2], 2)
q =
    1.999142573254893e+00    -1.058731468865185e-06
   -9.995712866274463e-01     5.293657344325927e-07
           0                2.132104856281811e-22
  -1.001419523209202e+00     5.293657344325930e-07
    2.002839046418403e+00    -1.058731468865186e-06
           0                0
    9.990107782605554e-01    -5.293657344325924e-07
    1.998021556521111e+00    -1.058731468865185e-06
           0                0
    1.000440759343655e+00    -2.804966567016011e-06
    2.000881518687309e+00    -5.609933134032022e-06
   -1.993817021879831e+00     5.609933134032020e-06
           0                0
    9.969085109399156e-01    -2.804966567016010e-06
   -2.005318289232261e+00     5.609933134032017e-06
    1.002659144616131e+00    -2.804966567016008e-06
```

Notice that the relative difference between \mathbf{p} and \mathbf{q} is of the order 10^{-3} . The vector `delta` is of norm approximately equal to $1.1055 \cdot 10^{-5}$. According to Lemma 5, the norm of this vector should be of the order $10^{-4} = \|10^{-2}\mathbf{v}\|^2$, which is indeed the case. In addition, the Iterated Scaling Lemma claims that `delta` should be perpendicular to the kernel of Terracini's matrix \mathbf{Tr} . Computing `norm(V(:, 13:16)'*delta)` one finds approximately $4.2999 \cdot 10^{-21}$, which is a fairly good approximation to zero. Finally, Lemma 5

FIGURE 1. Evolution of $\|\nabla^{(k)}\|$ in function of the iteration number, for different initial values of $\|\nabla^{(1)}\| = 10^{-q}$.



states that the vector \mathbf{q} should correspond with a different representative of the same rank-2 decomposition as represented by the vectorized factor matrices \mathbf{p} . We can test this property by computing the norm of the difference between the individual rank-1 tensors represented respectively by \mathbf{q} and \mathbf{p} :

```
norm( kron(q(1:3),kron(q(4:6),q(7:8))) - kron(A(:,1),kron(B(:,1),C(:,1))) )
ans =
    0
norm( kron(q(9:11),kron(q(12:14),q(15:16))) - kron(A(:,2),kron(B(:,2),C(:,2))) )
ans =
    1.241267076623637e-15
```

As a final experiment involving the Iterated Scaling Lemma, I present in Figure 1 the value of $\|\nabla^{(k)}\|$ in the proof of Lemma 5 for the perturbation vector $\mathbf{nabla} = 10^{-q} \cdot \mathbf{v}$ whereby $q = 1, 2, \dots, 5$. The figure illustrates the linear convergence claimed in (H1). Matlab's variable precision arithmetic capabilities using 100 digits of accuracy were employed for generating this figure.

9.2. The main theorem. For illustrating Theorem 7, I propose investigating the quantity on the left hand side of

$$(48) \quad \frac{\|\mathbf{p} - \hat{\mathbf{p}}\|/\|\mathbf{p}\|}{\|f(\mathbf{p}) - f(\hat{\mathbf{p}})\|/\|f(\mathbf{p})\|} \geq \frac{d(\mathbf{p}, \hat{\mathbf{p}})/\|\mathbf{p}\|}{\|f(\mathbf{p}) - f(\hat{\mathbf{p}})\|/\|f(\mathbf{p})\|}$$

as a proxy for the right hand side. We may estimate the condition number κ by generating a very large number of vectors $\hat{\mathbf{p}}$ such that $\|f(\mathbf{p}) - f(\hat{\mathbf{p}})\| \leq \varepsilon$, for some small value of ε , and then taking the maximum value. For instance, let us generate one random positive rank-2 decomposition in $\mathbb{R}^{3 \times 3 \times 2}$ as follows:

```
A = rand(3,2); B = rand(3,2); C = rand(2,2);
[c,F,~] = cpdcond({A,B,C});
```

In my experiment,

F{1} =	F{2} =	F{3} =
5.1518e-01 8.8821e-01	1.9032e-01 7.5082e-01	7.2302e-01 6.9447e-01
4.9802e-01 3.6941e-01	5.4218e-01 1.6653e-01	4.9879e-01 6.7487e-01
5.0806e-01 1.1117e-01	6.6436e-01 5.8845e-01	

and the relative norm-balanced condition number was $\kappa \approx 18.410308$. Then, the factor matrices \mathbf{F} were perturbed randomly as follows

```
Fp = arrayfun(...
    @(k) F{k}+1e-4*norm(F{k},'fro')*rand(size(F{k})), 1:3, 'UniformOutput', false ...
);
```

Starting from these factor matrices, I instructed the `cpd_nls` algorithm from Tensorlab to compute an approximate tensor rank decomposition of the tensor $f(\mathbf{p})$, where \mathbf{p} are the vectorized factor matrices corresponding to \mathbf{F} . Let $\hat{\mathbf{p}}$ denote the set of norm-balanced vectorized factor matrices obtained from a call to `cpd_nls`. The forgoing computation is repeated 10 million times; for every instance where $\|f(\mathbf{p}) - f(\hat{\mathbf{p}})\| \leq 10^{-14} \cdot \|f(\mathbf{p})\|$, the value on the left-hand side of (48) is computed. This resulted in 9,816,827 valid samples. Taking the maximum value over all the valid samples yielded approximately 10.437311 as estimate for the condition number. Given the very large number of trials, this is poor approximation to κ . In fact, the mean estimate of the condition number over all trials was only 0.8667234. These results seem to suggest that most perturbations—in this example—are actually very good for the tensor rank decomposition problem, because the parameters \mathbf{p} change relatively less than the tensor. However, if we apply the worst possible perturbation from Part III of the proof of Theorem 7 by executing the following code¹²

```
[~,F,Tr] = cpdcond({A,B,C});
[~,~,V] = svd(Tr); v = V(:,2*(3+3+2-2));
Fbad = {...
    F{1}+1e-8*[v(1:3) v(9:11)], ...
    F{2}+1e-8*[v(4:6) v(12:14)], ...
    F{3}+1e-8*[v(7:8) v(15:16)] ...
};
[~,Fbad,~] = cpdcond(Fbad);
p = flatten(F); pp = flatten(Fbad); T = cpdgen(F);
goodCondApp = (frob(pp-p)/frob(p)) / (frob(T - cpdgen(Fbad))/frob(T))
```

then we get a substantially more accurate estimate of the condition number, namely 18.410307—a relative difference of only about $5 \cdot 10^{-8}$ with respect to the true condition number.

9.3. The elementary properties. Let us numerically check Proposition 12. The condition numbers of rank-1 tensors of successively higher order behave as $\sqrt{d^{-1}}$, as evidenced by the next experiment:

```
a = orth(randn(2,1)); b = orth(randn(3,1)); c = orth(randn(4,1));
Conds = zeros(8,3);
for k = 3 : 10
    Conds(k-2,1) = cpdcond(arrayfun(@(x) 1*a, 1:k, 'UniformOutput', false));
    Conds(k-2,2) = cpdcond(arrayfun(@(x) 2*b, 1:k, 'UniformOutput', false));
    Conds(k-2,3) = cpdcond(arrayfun(@(x) 3*c, 1:k, 'UniformOutput', false));
end
norm( Conds.^(-2) - (3:10)'*ones(1,3) )
```

Matlab's output was $1.0272\text{e-}11$, illustrating that the relative condition number of a rank-1 tensor indeed only depends on the order of the tensor and, in particular, does not depend on the dimensions n_i of the tensor and its norm.

From the expression in Proposition 14 it is not difficult to derive that ill-conditioned weak 3-orthogonal tensors can be constructed in two principal ways, namely either by letting $\alpha_r \rightarrow 0$ while holding the other α_i 's constant, or by letting $\alpha_1 \rightarrow \infty$ while keeping the others constant. Observe that these two strategies are equivalent because the relative condition number is invariant under scaling. Consider for example the following weak 3-orthogonal (and odeco) tensors

$$\mathfrak{A}_s = 10^{-s} \mathbf{u}^1 \otimes \mathbf{u}^2 \otimes \mathbf{u}^3 + \mathbf{v}^1 \otimes \mathbf{v}^2 \otimes \mathbf{v}^3,$$

where $\langle \mathbf{u}^j, \mathbf{v}^j \rangle = 0$ for $j = 1, 2, 3$. The relative condition numbers for successive values of s are

s	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
$\log_{10}(\kappa(\mathfrak{A}_s))$	-0.2	0.4	1.1	1.8	2.4	3.1	3.8	4.4	5.1	5.8	6.4	7.1	7.8	8.4	9.1	9.8

Note that $\log_{10}(\kappa(\mathfrak{A}_s)) \approx -0.2 + 2s/3$ as can be understood by applying Proposition 14 to this example.

¹²The `flatten` function is also included in the ancillary files and essentially implements the `vecr` map.

9.4. An ill-conditioned example. Consider the following sequence of tensors

$$(49) \quad \mathfrak{A}_s = f(\mathbf{p}_s) = \mathbf{a}_1^1 \otimes \mathbf{a}_1^2 \otimes (\mathbf{x} + 2^{-s} \mathbf{a}_1^3) + \mathbf{a}_2^1 \otimes \mathbf{a}_2^2 \otimes (\mathbf{x} + 2^{-s} \mathbf{a}_2^3),$$

where $\mathbf{a}_i^k \in \mathbb{F}^{n_k}$ are random vectors sampled from a Gaussian distribution, $\mathbf{x} \in \mathbb{F}^{n_3}$ is any vector, and \mathbf{p}_s are the norm-balanced vectorized factor matrices. Every tensor \mathfrak{A}_s on this sequence is robustly 2-identifiable by Kruskal's theorem with probability 1. However,

$$\lim_{s \rightarrow \infty} \mathfrak{A}_s = (\mathbf{a}_1^1 \otimes \mathbf{a}_1^2 + \mathbf{a}_2^1 \otimes \mathbf{a}_2^2) \otimes \mathbf{x}$$

which is not a 2-identifiable tensor because

$$(\mathbf{x}_1^1 \otimes \mathbf{x}_1^2 + \mathbf{x}_2^1 \otimes \mathbf{x}_2^2) \otimes \mathbf{x} \quad \text{with} \quad \begin{bmatrix} \mathbf{x}_1^1 & \mathbf{x}_2^1 \end{bmatrix} = \begin{bmatrix} \mathbf{a}_1^1 & \mathbf{a}_2^1 \end{bmatrix} Z \quad \text{and} \quad \begin{bmatrix} \mathbf{x}_1^2 & \mathbf{x}_2^2 \end{bmatrix} = \begin{bmatrix} \mathbf{a}_1^2 & \mathbf{a}_2^2 \end{bmatrix} Z^{-1}$$

for any invertible $Z \in \mathbb{F}^{2 \times 2}$ is an alternative decomposition. Hence, \mathfrak{A}_∞ has infinitely many decompositions of length 2, yet it is also the limit of a sequence of 2-identifiable tensors.¹³

Because \mathfrak{A}_s is 2-identifiable and its rank is not larger than any of its multilinear ranks, it follows that a direct decomposition algorithm such as those in [32, 36, 37, 56] may be employed for computing \mathfrak{A}_s 's decomposition. Let $\hat{\mathbf{p}}_s \in f^\dagger(\mathfrak{A}_s)$ be the norm-balanced set of vectorized factor matrices that were computed by some numerical algorithm for solving the tensor rank decomposition problem (with rank equal to 2 in this case). Note that the numerical algorithm thus computed the decomposition of a nearby tensor $\hat{\mathfrak{A}}_s = f(\hat{\mathbf{p}}_s)$. The relative backward error is then defined as $\|\mathfrak{A}_s - \hat{\mathfrak{A}}_s\|/\|\mathfrak{A}_s\|$; it is a measure of how close the rank decomposition problem that was solved by the numerical algorithm was to the true rank decomposition problem. Whenever the backward error of an algorithm is always of the order of the machine precision, then the algorithm is called *backward stable*. As mentioned in the introduction, since one does not know the true solution in a practical setting, the relative forward error $d(\mathbf{p}_s, f^\dagger(\mathfrak{A}_s))\|/\|\mathbf{p}_s\|$ —which is I believe to be the quantity of interest in any data-analysis application where the individual rank-1 tensors appearing in the tensor rank decomposition are to be interpreted—cannot be evaluated. However, the condition number proposed in this paper allows us to (asymptotically) bound the relative forward error from above, because of the relation

$$(50) \quad \frac{d(\mathbf{p}_s, \hat{\mathbf{p}}_s)}{\|\mathbf{p}_s\|} \lesssim \kappa(\mathfrak{A}_s) \cdot \frac{\|\mathfrak{A}_s - \hat{\mathfrak{A}}_s\|}{\|\mathfrak{A}_s\|}.$$

In the present experiment, we will investigate these four quantities, i.e., the relative forward error, the condition number, the relative backward error, and the estimated upper bound on the right hand side in the last inequality.

As a concrete experiment, let us investigate $13 \times 11 \times 7$ tensors whose factor matrices $\mathbf{F}(\mathbf{s})$ were generated as follows:

```
A = randn(13,2); B = randn(11,2); C = randn(7,2); x = randn(7,1);
F = @(s) {A, B, x*[1 1] + 2^(-s) * C};
```


I will use the `cpd_gevd` algorithm from Tensorlab for computing a direct decomposition of \mathfrak{A}_s for the values $s = 1, 2, \dots, 45$. For simplicity, I will use $\|\mathbf{p}_s - \hat{\mathbf{p}}_s\|$ as a proxy for the forward error $d(\mathbf{p}_s, \hat{\mathbf{p}}_s)$; by definition, the former is an upper bound on the latter. The backward error, the proxy of the forward error, and estimated forward error obtained by multiplying the condition number with the backward error are all plotted in function of s in Figure 2.

I believe that Figure 2 clearly illustrates the danger of simply investigating the relative backward error. One sees that the backward error of the generalized eigenvalue decomposition algorithm is of the order of the machine precision so long as the tensor \mathfrak{A}_s is sufficiently¹⁴ discernible from \mathfrak{A}_∞ ; for the latter both the generalized eigendecomposition and the tensor rank decomposition no longer admit a unique solution. The forward error, however, may be several orders of magnitude larger. In this

¹³There is nothing special about the coefficient 2^{-s} in front of \mathbf{a}_i^3 ; one could equally well have replaced it with any decreasing function of s —even different coefficients in front of \mathbf{a}_1^3 and \mathbf{a}_2^3 would yield the same type of behavior, i.e., a sequence of 2-identifiable tensors converging to a tensor with an infinite number of decompositions.

¹⁴When $s = 45$, Tensorlab's implementation of the ST-HOSVD [79] algorithm detects that the multilinear rank of \mathfrak{A}_s is very close to $(2, 2, 1)$, at which point the software prudently refuses to employ the generalized eigenvalue decomposition algorithm.

FIGURE 2. The relative backward error, relative forward error, and relative condition number of a particular instance of the sequence in (49).



example, it can be observed that the forward error grows as s increases because of the increase of the condition number from roughly $\mathcal{O}(1)$ to $\mathcal{O}(10^{13})$.

This experiment seems to suggest that the condition number strongly increases as one moves closer to a tensor with ∞ many decompositions. An intuitive explanation is that an identifiable tensor \mathfrak{A}_ϵ , whose norm-balanced vectorized factor matrices are given by \mathbf{p}_ϵ , at relative distance ϵ to a tensor with ∞ many decompositions \mathfrak{A} needs but a perturbation of magnitude ϵ for *potentially* increasing the relative forward error to $\mathcal{O}(1)$. Then,

$$\kappa(\mathfrak{A}_\epsilon) \gtrsim \frac{d(\mathbf{p}_\epsilon, f^\dagger(\mathfrak{A})) / \|\mathbf{p}\|}{\|\mathfrak{A} - \mathfrak{A}_\epsilon\| / \|\mathfrak{A}_\epsilon\|} = \mathcal{O}(1) \cdot \frac{\|\mathfrak{A}_\epsilon\|}{\|\mathbf{p}_\epsilon\|} \cdot \epsilon^{-1},$$

which diverges to infinity if the fraction in the right hand side remains constant as $\epsilon \rightarrow 0$. Further research is necessary for proving this theoretically.

9.5. Two ill-posed examples. According to Demmel [34], a problem that is close to an ill-posed problem is often ill-conditioned. We will investigate numerically whether this property is also admitted by the proposed condition number, starting with the well-known example of de Silva and Lim [33]:

$$(51) \quad \mathfrak{A}_s = f(\mathbf{p}_s) = 2^{s/5}(\mathbf{a}^1 + 2^{-s/5}\mathbf{b}^1) \otimes (\mathbf{a}^2 + 2^{-s/5}\mathbf{b}^2) \otimes (\mathbf{a}^3 + 2^{-s/5}\mathbf{b}^3) - 2^{s/5}\mathbf{a}^1 \otimes \mathbf{a}^2 \otimes \mathbf{a}^3,$$

where $\mathbf{a}^k, \mathbf{b}^k \in \mathbb{F}^{n_k}$ are linearly independent vectors. Every tensor on this sequence is robustly 2-identifiable by Kruskal's criterion ($2 \leq \frac{1}{2}(2+2+2-2)$), and the limit of this sequence of rank-2 tensors is the rank-3 tensor


$$\lim_{s \rightarrow \infty} \mathfrak{A}_s = \mathbf{b}^1 \otimes \mathbf{a}^2 \otimes \mathbf{a}^3 + \mathbf{a}^1 \otimes \mathbf{b}^2 \otimes \mathbf{a}^3 + \mathbf{a}^1 \otimes \mathbf{a}^2 \otimes \mathbf{b}^3.$$

As in the foregoing ill-conditioned example, the tensor rank decomposition of \mathfrak{A}_s can be computed using a direct algorithm, such as `cpd_gevd`. Let $\hat{\mathbf{p}}_s$ denote the set of norm-balanced vectorized factor matrices obtained from applying this numerical algorithm to \mathfrak{A}_s , and let $\hat{\mathfrak{A}}_s = f(\hat{\mathbf{p}}_s)$.

As a particular instance, I generated vectors $\mathbf{a}^1, \mathbf{b}^1 \in \mathbb{R}^5$, $\mathbf{a}^2, \mathbf{b}^2 \in \mathbb{R}^4$ and $\mathbf{a}^3, \mathbf{b}^3 \in \mathbb{R}^3$ all of whose entries were sampled from a standard normal distribution. Then, for all $s = 5, 6, \dots, 100$, the rank-2 decomposition was computed with `cpd_gevd`, recording the relative backward error, the proxy of the relative forward error, and the relative condition number at $\hat{\mathfrak{A}}_s$ multiplied with the relative backward error. These quantities are plotted in Figure 3.

It is again clear from Figure 3 that the backward error is not a good measure of the stability of the vectorized factor matrices, as the forward error is several orders of magnitude larger for large s . Contrary to the ill-conditioned example from the foregoing subsection, the backward error is also increasing significantly. It turns out that the increase of the backward error is quite strongly correlated with the fraction $\|\hat{\mathbf{p}}_s\| / \|\hat{\mathfrak{A}}_s\|$ with a correlation coefficient of about 0.76.

FIGURE 3. The relative backward error, relative forward error, and relative condition number of a particular instance of the sequence in (51).



From the figure it can be deduced that the condition number increases from $\mathcal{O}(1)$ to $\mathcal{O}(10^{10})$ as s increases from 5 to 90, suggesting that the conditioning of the problem deteriorates as one moves closer to the ill-posed tensor rank decomposition problem. Notice that the upper bound on the forward error seems to stagnate around $s = 85$. It should be stressed that this is a numerical issue in the computation of the condition number. Employing variable precision arithmetic, it can be verified that the condition number keeps increasing as s increases. Fortunately, the occurrence of these numerical difficulties in computing the condition number can be detected by verifying that the N th singular value of Terracini's matrix is larger than a small constant multiple of the machine precision— $\epsilon_{\text{mach}} \approx 2.2 \cdot 10^{-16}$ using standard double precision in Matlab—multiplied with the largest singular value σ_1 of Terracini's matrix. The basic implementation that is included with this manuscript implements this test and issues a warning if the N th singular value is larger than $100 \cdot \epsilon_{\text{mach}} \cdot \sigma_1$. In this particular example, the code issues warnings starting from $s \geq 75$. Although I have no formal proof, I believe that it is reasonable to assume that when the computation of the condition number starts to suffer from numerical difficulties, the tensor rank decomposition problem is likely ill-conditioned.

It seems that the behavior of the condition number is very similar in other known examples of sequences tending to an ill-posed tensor rank decomposition problem. I present one more case, namely the example of Paatero [61, Section 7]:


```
Ab = randn(n1,3); Bb = randn(n2,3); Cb = randn(n3,3);
A = @(e) [-Ab(:,1)/e - Ab(:,2)/e, Ab(:,1)/e + e*e*Ab(:,3)/2, Ab(:,2)/e];
B = @(e) [-Bb(:,1)/e, Bb(:,1)/e + e*e*Bb(:,2)/2, Bb(:,1)/e + e*e*Bb(:,3)/2];
C = @(e) [-Cb(:,1)/e, Cb(:,1)/e + e*e*Cb(:,2)/2, Cb(:,1)/e + e*e*Cb(:,3)/2];
F = @(s) {A(2^(-s/16)), B(2^(-s/16)), C(2^(-s/16))};
```

wherein $n_1, n_2, n_3 \geq 3$. I chose $n_1 = 5$, $n_2 = 4$ and $n_3 = 3$ for the present experiment. Then, every tensor \mathfrak{A}_s whose factor matrices are given by $F(s)$ will be robustly r -identifiable by Kruskal's theorem ($3 \leq \frac{1}{2}(3 + 3 + 3 - 2)$), provided that the matrices Ab , Bb and Cb have linearly independent columns.

In Figure 4, the evolution of the relative backward error, the proxy of the relative forward error, and the condition number multiplied with the backward error are plotted in function of s . The behavior is very similar to the case of de Silva and Lim's example. The relative backward error is increasing from $\mathcal{O}(10^{-16})$ to $\mathcal{O}(10^{-10})$, seemingly driven by the increase of the fraction $\|\hat{\mathbf{p}}_s\|/\|\mathfrak{A}_s\|$ —the correlation coefficient is about 0.73 in this case. Meanwhile, the relative condition number increases from $\mathcal{O}(10^2)$ to $\mathcal{O}(10^6)$. It is again visually clear that the condition number multiplied with the backward error provides a more accurate estimate of the forward error than the backward error does. For $s = 100$, the code warned about a possibly inaccurate computation of the condition number. Also note that the `cpd_gevd` failed to compute a reasonable decomposition in this case, for some reason unbeknown to me.

These experiments seem to suggest that the condition number of the rank- r decomposition problem strongly increases as one moves closer to a tensor \mathfrak{A} that can be approximated arbitrarily well by a

FIGURE 4. The relative backward error, relative forward error, and relative condition number of a particular instance of Paatero’s example [61, Section 7].



rank- r tensor, but which does not admit an exact rank- r decomposition. It is known from [33] that on such a sequence, $\|\mathbf{p}_s\| \rightarrow \infty$, i.e., the norm of the vectorized factor matrices grows without bound. On the other hand, $\|\mathfrak{A}_s\|$ is bounded by a constant for sufficiently large s , because the sequence $\mathfrak{A}_s \rightarrow \mathfrak{A}$. Since

$$\kappa(\mathfrak{A}_\epsilon) \gtrsim \frac{d(\mathbf{p}_\epsilon, f^\dagger(\mathfrak{A})) / \|\mathbf{p}_\epsilon\|}{\|\mathfrak{A} - \mathfrak{A}_\epsilon\| / \|\mathfrak{A}_\epsilon\|} = d(\mathbf{p}_\epsilon, f^\dagger(\mathfrak{A})) \frac{1}{\|\mathbf{p}_\epsilon\| \cdot \|\mathfrak{A} - \mathfrak{A}_\epsilon\| / \|\mathfrak{A}_\epsilon\|},$$

no conclusions can be drawn without understanding the relative speed with which $\|\mathbf{p}_s\| \rightarrow \infty$ and $\|\mathfrak{A} - \mathfrak{A}_\epsilon\| / \|\mathfrak{A}_\epsilon\| \rightarrow 0$, and what happens to $d(\mathbf{p}_\epsilon, f^\dagger(\mathfrak{A}))$ —the latter is not even well-defined, because the fiber is empty. The two examples seem to suggest that the condition number becomes unbounded.

10. CONCLUSIONS

A condition number for the tensor rank decomposition problem was introduced. Provided that the input tensor satisfies a technical assumption called robust r -identifiability, the proposed condition number admits a particularly simple interpretation as the asymptotically worst fraction between the relative forward and relative backward error. The investigated condition number corresponds to a multiple of the inverse of a certain singular value of Terracini’s matrix. Several basic properties of the norm-balanced condition number were investigated. The main conclusions of this paper are that

- Terracini’s matrix is the key to investigating conditioning of robustly r -identifiable tensor rank decompositions;
- the condition number multiplied with the backward error yields an asymptotic upper bound on the forward error;
- rank-1 tensors are always well-conditioned;
- weak 3-orthogonal tensors (which include odeco tensors) are well-conditioned when the norms of the individual rank-1 terms are roughly of the same order of magnitude; and
- weak 3-orthogonal tensors are ill-conditioned when several orders of magnitude of difference exists in the norms of the rank-1 terms.

In addition, the numerical experiments provide some preliminary evidence suggesting that robustly r -identifiable tensors that are close to

- tensors with ∞ many rank- r decompositions (see Section 9.4), or
- tensors with border rank r but rank strictly larger (see Section 9.5)

can be ill-conditioned.

I hope that this paper opens several avenues of further research as well as practical application. In particular, I hope that it may find application in a more rigorous stability analysis of computed tensor rank decompositions. I believe that it may assist in understanding the convergence of gradient-based

optimization methods for computing approximate tensor rank decompositions. Another interesting question concerns the conditioning of the tensor rank *approximation* problem, whereby the sensitivity of the factor matrices is investigated with respect to a general, unstructured perturbation of a tensor. It is also an open question whether the behavior observed in Sections 9.4 and 9.5 may be explained theoretically. Finally, it may be anticipated that the analysis and techniques presented in this paper may also be applied in more structured variants of the tensor rank decomposition, such as Waring decompositions of symmetric tensors [29]—which are provably generically r -identifiable for all subgeneric symmetric ranks, and thus generically robustly r -identifiable, because of [27]—or of tensor rank decompositions of partially symmetric tensors.

REFERENCES

- [1] H. ABO, G. OTTAVIANI, AND C. PETERSON, *Induction for secant varieties of Segre varieties*, Trans. Amer. Math. Soc., 361 (2009), pp. 767–792.
- [2] E. ACAR, D.M. DUNLAVY, T.G. KOLDA, AND M. MØRUP, *Scalable tensor factorizations for incomplete data*, Chemometr. Intell. Lab., 106 (2011), pp. 41–56.
- [3] E. S. ALLMAN, C. MATIAS, AND J. A. RHODES, *Identifiability of parameters in latent structure models with many observed variables*, Ann. Stat., 37 (2009), pp. 3099–3132.
- [4] A. ANANDKUMAR, R. GE, D. HSU, S. M. KAKADE, AND M. TELGARSKY, *Tensor decompositions for learning latent variable models*, J. Mach. Learn. Res., 15 (2014), pp. 2773–2832.
- [5] C.J. APPELLOF AND E.R. DAVIDSON, *Strategies for analyzing data from video fluorometric monitoring of liquid chromatographic effluents*, Anal. Chem., 53 (1981), pp. 2053–2056.
- [6] A. BERNARDI, G. BLEKHERMAN, AND G. OTTAVIANI, *On real typical ranks*, arXiv:1512.01853, (2015).
- [7] A. BHASKARA, M. CHARIKAR, AND A. VIJAYARAGHAVAN, *Uniqueness of tensor decompositions with applications to polynomial identifiability*, JMLR: Workshop and Conference Proceedings, 35 (2014), pp. 1–37.
- [8] D. BINI, *Relations between exact and approximate bilinear algorithms*, Calcolo, 17 (1980), pp. 87–97.
- [9] D. BINI, G. LOTTI, AND F. ROMANI, *Approximate solutions for the bilinear form computational problem*, SIAM J. Comput., 9 (1980), pp. 692–697.
- [10] G. BLEKHERMAN, *Typical real ranks of binary forms*, Found. Comput. Math., 15 (2015), pp. 793–798.
- [11] C. BOCCI AND L. CHIANTINI, *On the identifiability of binary Segre products*, J. Algebraic Geom., 22 (2013), pp. 1–11. 10.1090/S1056-3911-2011-00592-4.
- [12] C. BOCCI, L. CHIANTINI, AND G. OTTAVIANI, *Refined methods for the identifiability of tensors*, Ann. Mat. Pura Appl. (4), (2013).
- [13] J. BOCHNAK, M. COSTE, AND M. ROY, *Real Algebraic Geometry*, Springer-Verlag, 1998.
- [14] A. BORALEVI, J. DRAISMA, E. HOROBET, AND E. ROBEVA, *Orthogonal and unitary tensor decomposition from an algebraic perspective*, arXiv:1512.08031, (2015).
- [15] P. BÜRGISSER AND F. CUCKER, *Condition: The Geometry of Numerical Algorithms*, vol. 349 of Grundlehren der mathematischen Wissenschaften, Springer-Verlag, 2013.
- [16] J. CARROLL AND J.-J. CHANG, *Analysis of individual differences in multidimensional scaling via an n -way generalization of “Eckart–Young” decomposition*, Psychometrika, 35 (1970), pp. 283–319.
- [17] M.V. CATALISANO, A.V. GERAMITA, AND A. GIMIGLIANO, *Ranks of tensors, secant varieties of Segre varieties and fat points*, Linear Algebra Appl., 355 (2002), pp. 263–285.
- [18] ———, *Higher secant varieties of the Segre varieties $\mathbb{P}^1 \times \dots \times \mathbb{P}^1$* , J. Pure Appl. Algebra, 201 (2005), pp. 367–380.
- [19] ———, *Segre–Veronese embeddings of $\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1$ and their secant varieties*, Collect. Math., 1 (2007), pp. 1–24.
- [20] ———, *On the ideals of secant varieties to certain rational varieties*, J. Algebra, 319 (2008), pp. 1913–1931.
- [21] ———, *Secant varieties of $\mathbb{P}^1 \times \dots \times \mathbb{P}^1$ (n -times) are not defective for $n \geq 5$* , J. Algebraic Geom., 20 (2011), pp. 295–327.
- [22] J. CHEN AND Y. SAAD, *On the tensor SVD and the optimal low rank orthogonal approximation of tensors*, SIAM J. Matrix Anal. Appl., 30 (2009), pp. 1709–1734.
- [23] L. CHIANTINI AND C. CILIBERTO, *Weakly defective varieties*, Trans. Amer. Math. Soc., 354 (2001), pp. 151–178.
- [24] ———, *On the concept of k -secant order of a variety*, J. London Math. Soc. (2), 73 (2006), pp. 436–454.
- [25] L. CHIANTINI AND G. OTTAVIANI, *On generic identifiability of 3-tensors of small rank*, SIAM J. Matrix Anal. Appl., 33 (2012), pp. 1018–1037.
- [26] L. CHIANTINI, G. OTTAVIANI, AND N. VANNIEUWENHOVEN, *An algorithm for generic and low-rank specific identifiability of complex tensors*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 1265–1287.
- [27] ———, *On generic identifiability of symmetric tensors*, Trans. Amer. Math. Soc., (2016). Accepted.
- [28] P. COMON, *Independent component analysis, a new concept?*, Signal Proc., 36 (1994), pp. 287–314.
- [29] P. COMON, G. H. GOLUB, L.-H. LIM, AND B. MOURRAIN, *Symmetric tensors and symmetric tensor rank*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 1254–1279.
- [30] P. COMON AND G. OTTAVIANI, *On the typical rank of real binary forms*, Linear Multilinear Algebra, 60 (2012), pp. 657–667.

- [31] P. COMON, J.M.F. TEN BERGE, L. DE LATHAUWER, AND J. CASTAING, *Generic and typical ranks of multi-way arrays*, Linear Algebra Appl., 430 (2009), pp. 2997–3007.
- [32] L. DE LATHAUWER, *A link between the canonical decomposition in multilinear algebra and simultaneous matrix diagonalization*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 642–666.
- [33] V. DE SILVA AND L.-H. LIM, *Tensor rank and the ill-posedness of the best low-rank approximation problem*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 1084–1127.
- [34] J. W. DEMMEL, *The geometry of ill-conditioning*, J. Complexity, 3 (1987), pp. 201–229.
- [35] I. DOMANOV AND L. DE LATHAUWER, *On the uniqueness of the canonical polyadic decomposition of third-order tensors—part II: Uniqueness of the overall decomposition*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 876–903.
- [36] ———, *Canonical polyadic decomposition of third-order tensors: reduction to generalized eigenvalue decomposition*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 636–660.
- [37] ———, *Canonical polyadic decomposition of third-order tensors: relaxed uniqueness conditions and algebraic algorithm*, arXiv:1501.07251, (2015).
- [38] ———, *Generic uniqueness conditions for the canonical polyadic decomposition and INDSCAL*, SIAM J. Matrix Anal. Appl., 36 (2015), pp. 1567–1589.
- [39] A. L. DONTCHEV AND R. T. ROCKAFELLAR, *Implicit Function and Solution Mappings: A View from Variational Analysis*, Springer Monographs in Mathematics, Springer, 2009.
- [40] F. GESMUNDO, *An asymptotic bound for secant varieties of Segre varieties*, Ann. Univ. Ferrara, 59 (2013), pp. 285–302.
- [41] P. GRIFFITHS AND J. HARRIS, *Principles of Algebraic Geometry*, John Wiley & Sons, Inc., 1978.
- [42] J. HARRIS, *Algebraic Geometry, A First Course*, vol. 133 of Graduate Text in Mathematics, Springer-Verlag, 1992.
- [43] R. A. HARSHMAN, *Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multi-modal factor analysis*, UCLA Working Papers in Phonetics, 16 (1970), pp. 1–84.
- [44] J. D. HAUSTEIN, L. OEDING, G. OTTAVIANI, AND A. J. SOMMESE, *Homotopy techniques for tensor decomposition and perfect identifiability*, arXiv:1501.00090, (2015).
- [45] C. HAYASHI AND F. HAYASHI, *A new algorithm to solve PARAFAC-model*, Behaviormetrika, 11 (1982), pp. 49–60.
- [46] N. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, PA, 1996.
- [47] F.L. HITCHCOCK, *The expression of a tensor or a polyadic as a sum of products*, J. Math. Phys., 6 (1927), pp. 164–189.
- [48] ———, *Multiple invariants and generalized rank of a p-way matrix or tensor*, J. Math. Phys., 7 (1927), pp. 39–79.
- [49] R.A. HORN AND C.R. JOHNSON, *Matrix Analysis*, Cambridge University Press, New York, NY, USA, 1986.
- [50] T. JIANG AND N.D. SIDIROPOULOS, *Kruskal’s permutation lemma and the identification of CANDECOMP-/PARAFAC and bilinear models with constant modulus constraints*, IEEE Trans. Signal Process., 52 (2004), pp. 2625–2636.
- [51] T. G. KOLDA, *Orthogonal tensor decompositions*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 243–255.
- [52] Z. KOLDOVSKÝ, P. TICHAVSKÝ, AND A.-H. PHAN, *Stability analysis and fast damped Gauss–Newton algorithm for INDSCAL tensor decomposition*, in Proc. IEEE Workshop Statist. Signal Process., 2011, pp. 581–584.
- [53] J. B. KRUSKAL, *Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics*, Linear Algebra Appl., 18 (1977), pp. 95–138.
- [54] J.M. LANDSBERG, *Tensors: Geometry and Applications*, vol. 128 of Graduate Studies in Mathematics, American Mathematical Society, Providence, Rhode Island, 2012.
- [55] J. M. LEE, *Introduction to Smooth Manifolds*, vol. 218 of Graduate Texts in Mathematics, Springer, New York, USA, 2 ed., 2013.
- [56] S. E. LEURGANS, R. T. ROSS, AND R. B. ABEL, *A decomposition for three-way arrays*, SIAM J. Matrix Anal. Appl., 14 (1993), pp. 1064–1083.
- [57] X. LIU AND N. D. SIDIROPOULOS, *Cramér–Rao lower bounds for low-rank decomposition of multidimensional arrays*, IEEE Trans. Signal Process., 49 (2001), pp. 2074–2086.
- [58] I.V. OSELEDETS AND D.V. SAVOST’YANOV, *Minimization methods for approximating tensors and their comparison*, Comp. Math. Math. Phys., 46 (2006), pp. 1641–1650.
- [59] P. PAATERO, *A weighted non-negative least squares algorithm for three-way ‘PARAFAC’ factor analysis*, Chemometr. Intell. Lab., 38 (1997), pp. 223–242.
- [60] ———, *The Multilinear Engine—a table-driven, least squares program for solving multilinear problems, including the n-way parallel factor analysis model*, J. Comput. Graph. Stat., 8 (1999), pp. 854–888.
- [61] ———, *Construction and analysis of degenerate PARAFAC models*, J. Chemometr., 14 (2000), pp. 285–299.
- [62] A.-H. PHAN, P. TICHAVSKÝ, AND A. CICHOCKI, *Fast alternating LS algorithms for high order CANDECOMP-/PARAFAC tensor factorizations*, IEEE Trans. Signal Process., 61 (2013), pp. 4834–4846.
- [63] ———, *Low complexity damped Gauss–Newton algorithms for CANDECOMP/PARAFAC*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 126–147.
- [64] E. ROBEVA, *Orthogonal decomposition of symmetric tensors*, SIAM J. Matrix Anal. Appl., 37 (2016), pp. 86–102.
- [65] J. J. ROTMAN, *An Introduction to the Theory of Groups*, vol. 148 of Graduate Texts in Mathematics, Springer-Verlag, New York, USA, 4 ed., 1994.
- [66] N. D. SIDIROPOULOS AND R. BRO, *On the uniqueness of multilinear decomposition of N-way arrays*, J. Chemom., 14 (2000), pp. 229–239.

- [67] L. SORBER, M. VAN BAREL, AND L. DE LATHAUWER, *Optimization-based algorithms for tensor decompositions: canonical polyadic decomposition, decomposition in rank- $(l_r, l_r, 1)$ terms, and a new generalization*, SIAM J. Optim., 23 (2013), pp. 695–720.
- [68] ———, *Tensorlab v2.0*. <http://www.tensorlab.net>, 2014.
- [69] G. W. STEWART AND J.-G. SUN, *Matrix Perturbation Theory*, vol. 33 of Computer Science and Scientific Computing, Academic Press, 1990.
- [70] V. STRASSEN, *Rank and optimal computation of generic tensors*, Linear Algebra Appl., 52–53 (1983), pp. 645–685.
- [71] A. TERRACINI, *Sulla V_k per cui la varietà degli S_h $h + 1$ -secanti ha dimensione minore dell'ordinario*, Rend. Circ. Mat. Palermo (2), 31 (1911), pp. 392–396.
- [72] P. TICHAVSKÝ AND Z. KOLDOVSKÝ, *Weight adjusted tensor method for blind separation of underdetermined mixtures of nonstationary sources*, IEEE Trans. Signal Process., 59 (2011), pp. 1037–1047.
- [73] P. TICHAVSKÝ, A.-H. PHAN, AND Z. KOLDOVSKÝ, *Cramér–Rao-induced bounds for CANDECOMP/PARAFAC tensor decomposition*, IEEE Trans. Signal Process., 61 (2013), pp. 1986–1997.
- [74] G. TOMASI, *Practical and computational aspects in chemometric data analysis*, PhD thesis, The Royal Veterinary and Agricultural University, Frederiksberg, Denmark, 2006.
- [75] G. TOMASI AND R. BRO, *PARAFAC and missing values*, Chemometr. Intell. Lab., 75 (2005), pp. 163–180.
- [76] ———, *A comparison of algorithms for fitting the PARAFAC model*, Comput. Stat. Data An., 50 (2006), pp. 1700–1734.
- [77] N. VANNIEUWENHOVEN, K. MEERBERGEN, AND R. VANDEBRIL, *Computing the gradient in optimization algorithms for the CP decomposition in constant memory through tensor blocking*, SIAM J. Sci. Comput., 37 (2015), pp. C415–C438.
- [78] N. VANNIEUWENHOVEN, J. NICAISE, R. VANDEBRIL, AND K. MEERBERGEN, *On generic nonexistence of the Schmidt–Eckart–Young decomposition for complex tensors*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 886–903.
- [79] N. VANNIEUWENHOVEN, R. VANDEBRIL, AND K. MEERBERGEN, *A new truncation strategy for the higher-order singular value decomposition*, SIAM J. Sci. Comput., 34 (2012), pp. A1027–A1052.
- [80] ———, *A randomized algorithm for testing nonsingularity of structured matrices with an application to asserting nondefectivity of segre varieties*, IMA J. Numer. Anal., 35 (2015), pp. 289–324.
- [81] J. H. WILKINSON, *The Algebraic Eigenvalue Problem*, Oxford University Press, Amen House, London E.C. 4, United Kingdom, 1965.
- [82] T. ZHANG AND G. H. GOLUB, *Rank-one approximation to high order tensors*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 534–550.

APPENDIX A. MATLAB/OCTAVE PROGRAM CODE

```

function [ kappa, A, T ] = cpdcond( A, ~ )
%CPDCOND Computes the condition number.
% The routine estimates the norm-balanced condition number of the factor
% matrices A. If the second input argument is specified, the absolute
% condition number is computed.
% The first output argument is the norm-balanced condition number.
% The second output argument are the norm-balanced factor matrices with
% respect to which the condition number is computed.
% The third output argument contains Terracini's matrix of the
% norm-balanced representatives.

% Constants
d = length(A);
n = arrayfun(@(k) size(A{k},1), 1:d);
r = size(A{1},2);
Pi = prod(n);
Sigma = sum(n-1);

%% Apply norm-balancing
scal = ones(1,r);
for k = 1 : d
    lscal = arrayfun(@(l) norm(A{k}(:,l)), 1:r);
    scal = scal .* lscal.^(1/d);
    A{k} = A{k} * diag(1./lscal);
end
for k = 1 : d
    A{k} = A{k} * diag(scal);
end
normP = sqrt(d)*norm(scal);

% Condition number is unbounded by definition.
if r*(1+Sigma) > Pi
    warning('The number of terms is not subgeneric.')
    kappa = inf;
    T = [];
    return;
end

%% Construct Terracini's matrix
T = zeros(Pi, r*(d+Sigma));
for l = 1 : r
    lT = [];
    for k = 1 : d
        B = arrayfun(...
            @(kk) repmat(A{kk}(:,l),1,n(k)), 1:d, 'UniformOutput', false ...
        );
        B{k} = eye(n(k));
        lT = [lT krp(B)];
    end
    T(:,(d+Sigma)*(l-1)+1:(d+Sigma)*l) = lT;
end

%% Compute relative condition number
N = r*(1+Sigma);
S = svd(T,0);

```

```

kappa = 1 / S(N);
if nargin == 1
    tensNorm = norm(krp(A)*ones(r,1));
    kappa = (kappa * tensNorm) / normP;
end
if ( S(1)*100*eps > S(N) )
    warning('Computed condition number may not be accurate!')
end

% Compute Khatri-Rao product
function [X] = krp(B)
    Ns = prod(arrayfun(@(k) size(B{k},1),1:length(B)));
    X = zeros(Ns,size(B{1},2));
    for ll = 1 : size(X,2)
        x = B{1}(:,ll);
        for kk = 2 : length(B)
            x = kron(x,B{kk}(:,ll));
        end
        X(:,ll) = x;
    end
end

function [ p, delta ] = isl( p, nabla, n, r )
%ISL Applies the Iterated Scaling algorithm.
% The first input argument is the set of vectorized factor matrices.
% The second input argument is the perturbation vector.
% The third input argument is an array of integers containing the
% dimensions of the tensor represented by the vectorized factor
% matrices in p.
% The fourth input argument is the number of terms in the rank
% decomposition of the tensor represented by p.
% The first output argument is the rescaled set of vectorized factor
% matrices, representing the same tensor as p.
% The second output argument is a vector perpendicular to the kernel
% of Terracini's matrix, whose norm is approximately the square of the
% norm of nabla.

% Constants
d = length(n);
cn = cumsum(n);
Sigma = cn(d);

% Construct kernel
K = zeros(r*Sigma, r*(d-1));
li = 1;
for i = 1 : r
    pt = p((i-1)*Sigma+1:i*Sigma);
    for k = 2 : d
        off = (i-1)*Sigma;
        K(off+1:off+cn(1),li) = pt(1:cn(1));
        K(off+cn(k-1)+1:off+cn(k),li) = -pt(cn(k-1)+1:cn(k));
        li = li + 1;
    end
end

% Iterated scaling algorithm
p0 = p;

```

```

pk = p;
delta = zeros(size(p));
gam = ones(d-1,r);
[U,S,V] = svd(K,0);
nablanorm = inf;
while nablanorm > 10*eps
    v = V*(S\((U'*nabla)));
    gam = gam - reshape(v, [d-1 r]);
    zk = pk + K*v;
    off = 0;
    for i = 1 : r
        pk(off+1:off+cn(1)) = p0(off+1:off+cn(1))/prod(gam(:,i));
        off = off + n(1);
        for k = 2 : d
            pk(off+1:off+n(k)) = p0(off+1:off+n(k))*gam(k-1,i);
            off = off + n(k);
        end
    end
    nabla = zk - pk;
    delta = delta + (nabla - U*(U'*nabla));
    nabla = U*(U'*nabla);
    nablanorm = norm(nabla);
end
p = pk;
end

function [ v ] = flatten( A )
%FLATTEN Vectorizes a given set of factor matrices A.

v = [];
for r = 1 : size(A{1},2)
    for k = 1 : length(A);
        v = [v; A{k}(:,r)];
    end
end
end
end

```

KU LEUVEN, DEPARTMENT OF COMPUTER SCIENCE, B-3001 LEUVEN, BELGIUM
E-mail address: `nick.vannieuwenhoven@cs.kuleuven.be`